

Group Regularized Estimation under Structural Hierarchy

Yiyuan She, Zhifeng Wang and He Jiang

Department of Statistics, Florida State University

Abstract

Variable selection for models including interactions between explanatory variables often needs to obey certain hierarchical constraints. Weak or strong structural hierarchy requires that the existence of an interaction term implies at least one or both associated main effects to be present in the model. Lately this problem has attracted a lot of attention, but existing computational algorithms converge slow even with a moderate number of predictors. Moreover, in contrast to the rich literature on ordinary variable selection, there is a lack of statistical theory to show reasonably low error rates of hierarchical variable selection. This work investigates a new class of estimators that make use of multiple group penalties to capture structural parsimony. We show that the proposed estimators enjoy sharp rate oracle inequalities, and give the minimax lower bounds in strong and weak hierarchical variable selection. A general-purpose algorithm is developed with guaranteed convergence and global optimality. Simulations and real data experiments demonstrate the efficiency and efficacy of the proposed approach.

1 Introduction

In statistical applications, it is often noticed that an additive model including main effects only is inadequate. Including some higher-order terms, such as interactions, in particular, are often of great help in prediction and modeling.

Sometimes, interactions may be of independent interest; one example is the moderation analysis in behavioral sciences (Cohen et al., 2013). In this paper, we focus on the full quadratic model with all two-term interactions taken into account.

Let $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_p] \in \mathbb{R}^{n \times p}$ be the (raw) predictor matrix and $\mathbf{y} \in \mathbb{R}^n$ be the response vector. We assume the following nonlinear additive regression model

$$\mathbf{y} = b_0^* \mathbf{1} + \sum_{1 \leq j \leq p} b_j^* \mathbf{x}_j + \sum_{1 \leq j, k \leq p} \phi_{jk}^* \mathbf{x}_j \odot \mathbf{x}_k + \boldsymbol{\varepsilon}, \quad (1)$$

where $\boldsymbol{\varepsilon} \sim N(\mathbf{0}, \sigma^2 \mathbf{I})$ and \odot denotes the Hadamard product. The number of predictors is then $p + \binom{p}{2}$, posing a challenge in variable selection even when p is moderate. Moreover, in this scenario, statisticians are often interested in obtaining a model satisfying certain logical relations, such as the *structural hierarchy* discussed in Nelder (1977), McCullagh and Nelder (1989), and Hamada and Wu (1992). Hierarchy is a natural requirement in gene regulatory network studies (Davidson and Erwin, 2006), banded covariance matrix estimation (Bien et al., 2016) and lagged variable selection in time series. Hierarchical variable selection leads to reduced number of variables in measurement, referred to as *practical sparsity* (Bien et al., 2013). For instance, a model consisting of \mathbf{x}_1 , \mathbf{x}_2 , and $\mathbf{x}_1 \mathbf{x}_2$ may be more parsimonious to practitioners than a model involving \mathbf{x}_1 , \mathbf{x}_2 , and $\mathbf{x}_3 \mathbf{x}_4$. In our setting, there are two types of hierarchy (Chipman, 1996; Bien et al., 2013): strong hierarchy (**SH**) and weak hierarchy (**WH**). Let ϕ_{jk} be the coefficient of $\mathbf{x}_j \odot \mathbf{x}_k$ and $\mathbf{x}_k \odot \mathbf{x}_j$. SH means that if an interaction term exists in the model, then both of its associated main effects must be present, i.e., $\phi_{jk} \neq 0 \rightarrow b_j \neq 0$ and $b_k \neq 0$, while WH requires that the inclusion of an interaction implies at least one of its associated main effects to be added into the model, i.e., $\phi_{jk} \neq 0 \rightarrow b_j \neq 0$ or $b_k \neq 0$. We will show in Section 3 that WH is relatively easy to realize compared with SH. SH is invariant to linear transformations of predictors (Peixoto, 1990) and is the primary concern in this work.

It is a nontrivial task to maintain hierarchy in model selection using conventional approaches. LASSO (Tibshirani, 1996) may violate SH and WH as well. We refer to Nelder (1977), Peixoto (1987), Bickel et al. (2010), Wu et al. (2010), and Hao and Zhang (2014) for some well-developed multi-step procedures which, however, might be ad-hoc and greedy. This paper focuses

on *regularization-based* approaches. The past works in this direction include SHIM (Choi et al., 2010), VANISH (Radchenko and James, 2010) and HL (Bien et al., 2013). SHIM reparametrizes ϕ_{jk} as $\rho_{jk}b_jb_k$ and enforces sparsity in both $\mathbf{b} = [b_j]$ and $\boldsymbol{\rho} = [\rho_{jk}]$. The formulation is motivating, and we could also use $\phi_{jk} = \rho_{jk}(b_j^2 + b_k^2)$ for WH. However, the corresponding optimization problem is nonconvex and the computational algorithm of SHIM is quite slow in large-scale problems. VANISH is one of the main motivations of our work and will be discussed in detail in Section 3. HL is a recent breakthrough in hierarchical variable selection. One of its key ideas is to enforce a magnitude constraint on the coefficients, $\|\boldsymbol{\phi}_j\|_1 \leq |b_j|$, to make hierarchy naturally hold. Here, $\boldsymbol{\phi}_j$ is a vector of coefficients of the predictors $\mathbf{x}_j \odot \mathbf{x}_k$, $1 \leq k \leq p$. To handle the nonconvex constraint, Bien et al. (2013) rephrased it with the pseudo-positive and pseudo-negative parts b_j^+ , b_j^- of b_j but dropped all zero-product constraints $b_j^+b_j^- = 0$, $1 \leq j \leq p$. The quality of such a convex relaxation seems to have no theoretical justification in the literature. In our experience, HL has excellent performance when the main effects are strong and p is not very large. But it can miss some interaction effects and become computationally prohibitive on large datasets. For example, when $p = 1000$, HL can take days to obtain a 20-point solution path.

In this work, we propose and study group regularized estimation under structural hierarchy (**GRESH**). In theory, we are able to establish non-asymptotic oracle inequalities to show the error rates of the proposed estimators are minimax optimal up to some logarithm factors. We come up with a new recipe to conquer the theoretical difficulties when analyzing overlapping regularization terms in pursuing structural parsimony. Moreover, we develop a computational algorithm which guarantees the convergences of iterates and function values; it is not only efficient but also simple to implement.

The rest of the paper is organized as follows. Some notation and symbols are introduced in Section 2. Section 3 presents the general framework of GRESH. A fast computational algorithm with theoretical support is given in Section 4. Section 5 builds oracle inequalities for GRESH, and Section 6 shows the minimax optimal rates. In Section 7, simulation studies and real data analysis are conducted to show the prediction accuracy and computational efficiency of the proposed approach. All technical proofs, together with additional data analysis are given in the Appendices.

2 Notation

We introduce some convenient notation and symbols to be used throughout the paper. First, for any matrix $\mathbf{A} = [\mathbf{a}_1, \dots, \mathbf{a}_p]^\top \in \mathbb{R}^{p \times n}$, define its $(2, 1)$ -norm, $(2, \infty)$ -norm and ℓ_1 -norm as $\|\mathbf{A}\|_{2,1} = \sum_{i=1}^p \|\mathbf{a}_i\|_2$, $\|\mathbf{A}\|_{2,\infty} = \max_{1 \leq i \leq p} \|\mathbf{a}_i\|_2$ and $\|\mathbf{A}\|_1 = \|\text{vec}(\mathbf{A})\|_1$, respectively, where vec is the standard vectorization operator. The spectral norm and Frobenius norm of \mathbf{A} are denoted by $\|\mathbf{A}\|_2$ and $\|\mathbf{A}\|_F$, respectively. For any p -dimensional vector \mathbf{a} that is divided into K groups with \mathbf{a}_j representing the j -th subvector, its $(2, 1)$ -norm is defined by $\|\mathbf{a}\|_{2,1} = \sum_{j=1}^K \|\mathbf{a}_j\|_2$.

The following two operators diag and dg are introduced for notational simplicity. For a square matrix $\mathbf{A} := [a_{ij}]_{n \times n}$, $\text{diag}(\mathbf{A}) := [a_{11}, \dots, a_{nn}]^\top$, and for a vector $\mathbf{a} = [a_1, \dots, a_n]^\top \in \mathbb{R}^n$, $\text{diag}\{\mathbf{a}\}$ is defined as an $n \times n$ diagonal matrix with diagonal entries given by a_1, \dots, a_n . Define $\text{dg}(\mathbf{A}) := \text{diag}\{\text{diag}(\mathbf{A})\} = \text{diag}\{[a_{11}, \dots, a_{nn}]^\top\}$. We use $\mathbf{A}[\mathcal{I}, \mathcal{J}]$ to denote a submatrix of \mathbf{A} with rows and columns indexed by \mathcal{I} and \mathcal{J} , respectively.

For any arbitrary $\mathbf{b} \in \mathbb{R}^p$, $\Phi \in \mathbb{R}^{p \times p}$, we define

$$\begin{aligned} \mathcal{J}^{11}(\mathbf{b}, \Phi) &:= \{j \in [p] : b_j \neq 0, \phi_j \neq \mathbf{0}\}, \\ \mathcal{J}^{10}(\mathbf{b}, \Phi) &:= \{j \in [p] : b_j \neq 0, \phi_j = \mathbf{0}\}, \\ \mathcal{J}^{01}(\mathbf{b}, \Phi) &:= \{j \in [p] : b_j = 0, \phi_j \neq \mathbf{0}\}, \\ \mathcal{J}^{00}(\mathbf{b}, \Phi) &:= \{j \in [p] : b_j = 0, \phi_j = \mathbf{0}\}, \\ \mathcal{J}_e(\Phi) &:= \{j \in [p^2] : (\text{vec}(\Phi))_j \neq 0\}, \\ \mathcal{J}_G(\mathbf{b}, \Phi) &:= \{j \in [p] : b_j^2 + \|\phi_j\|_2^2 \neq 0\}, \end{aligned} \tag{2}$$

where $[p] = \{1, \dots, p\}$, and the coefficient vector ϕ_j denotes the j -th column of Φ . With $|\cdot|$ standing for set cardinality, we define $J^{11}(\mathbf{b}, \Phi) := |\mathcal{J}^{11}(\mathbf{b}, \Phi)|$, $J^{10}(\mathbf{b}, \Phi) := |\mathcal{J}^{10}(\mathbf{b}, \Phi)|$, $J^{01}(\mathbf{b}, \Phi) := |\mathcal{J}^{01}(\mathbf{b}, \Phi)|$, $J^{00}(\mathbf{b}, \Phi) := |\mathcal{J}^{00}(\mathbf{b}, \Phi)|$, $J_e(\Phi) := |\mathcal{J}_e(\Phi)|$ and $J_G(\mathbf{b}, \Phi) := |\mathcal{J}_G(\mathbf{b}, \Phi)|$. Clearly, $J^{11} + J^{10} + J^{01} + J^{00} = p$, and $J^{11} + J^{10} + J^{01} = J_G$. In addition, under SH, $J_G(\mathbf{b}, \Phi)$ equals the number of nonzero elements of \mathbf{b} . Given the true signal (\mathbf{b}^*, Φ^*) , the following abbreviated symbols are used: $J^{11*} = J^{11}(\mathbf{b}^*, \Phi^*)$, $J^{10*} = J^{10}(\mathbf{b}^*, \Phi^*)$, $J^{01*} = J^{01}(\mathbf{b}^*, \Phi^*)$, $J_e^* = J_e(\Phi^*)$, and $J_G^* = J_G(\mathbf{b}^*, \Phi^*)$.

In the paper, we frequently use the concatenated coefficient matrix for convenience

$$\Omega = [\mathbf{b}, \Phi^\top]^\top, \tag{3}$$

and its j -th column is denoted by $\boldsymbol{\Omega}_j = [b_j, \boldsymbol{\phi}_j^\top]^\top$. Given $\boldsymbol{\Omega}$, $\boldsymbol{\Omega}_b$ and $\boldsymbol{\Omega}_\Phi$ stand for $(\boldsymbol{\Omega}[1, :])^\top$ and $\boldsymbol{\Omega}[2 : (p+1), :]$, respectively.

Let $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_p] \in \mathbb{R}^{n \times p}$ be the raw predictor matrix. Define

$$\bar{\mathbf{X}} = [\mathbf{x}_1 \odot \mathbf{x}_1, \dots, \mathbf{x}_1 \odot \mathbf{x}_p, \dots, \mathbf{x}_p \odot \mathbf{x}_1, \dots, \mathbf{x}_p \odot \mathbf{x}_p] \in \mathbb{R}^{n \times p^2}, \quad (4)$$

$$\check{\mathbf{X}} = [\mathbf{x}_1, \mathbf{x}_1 \odot \mathbf{x}_1, \dots, \mathbf{x}_1 \odot \mathbf{x}_p, \dots, \mathbf{x}_p, \mathbf{x}_p \odot \mathbf{x}_1, \dots, \mathbf{x}_p \odot \mathbf{x}_p] \in \mathbb{R}^{n \times (p^2+p)}. \quad (5)$$

Then $\bar{\mathbf{X}}$ consists of all interactions, and $\check{\mathbf{X}}$ includes all p^2+p predictors in the quadratic model. Given any subset $\mathcal{J} \subset [p]$, we abbreviate $\mathbf{X}[:, \mathcal{J}]$ as $\mathbf{X}_{\mathcal{J}}$. It is also easy to see that $\text{diag}(\mathbf{X}\boldsymbol{\Phi}\mathbf{X}^\top) = \bar{\mathbf{X}} \text{vec}(\boldsymbol{\Phi})$ and $\mathbf{X}\mathbf{b} + \text{diag}(\mathbf{X}\boldsymbol{\Phi}\mathbf{X}^\top) = \check{\mathbf{X}} \text{vec}(\boldsymbol{\Omega})$.

For any two real numbers a and b , $a \lesssim b$ means that $a \leq b$ holds up to a multiplicative numerical constant. For two equally sized matrices $\mathbf{A} = (a_{ij})$ and $\mathbf{B} = (b_{ij})$, $\mathbf{A} \geq \mathbf{B}$ means $a_{ij} \geq b_{ij}$, for $\forall i, j$.

3 Group regularized estimation under structural hierarchy

For simplicity, we assume for now that there exists no intercept term in the model. Then (1) can be written as

$$\mathbf{y} = \mathbf{X}\mathbf{b}^* + \text{diag}(\mathbf{X}\boldsymbol{\Phi}^*\mathbf{X}^\top) + \boldsymbol{\varepsilon}, \quad (6)$$

where $\boldsymbol{\varepsilon} \sim N(\mathbf{0}, \sigma^2 \mathbf{I})$ with $\sigma^2 > 0$, $\mathbf{y} = [y_1, \dots, y_n]^\top \in \mathbb{R}^n$ is the response vector, and $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_p] \in \mathbb{R}^{n \times p}$ is the design matrix consisting of main effects only.

We describe a general framework for hierarchical variable selection, referred to as group regularized estimation under structural hierarchy or **GRESH**. GRESH has two different types, depending on which objects to regularize. Denoting the squared error loss by ℓ , i.e.,

$$\ell(\mathbf{b}, \boldsymbol{\Phi}) = \frac{1}{2} \|\mathbf{y} - \mathbf{X}\mathbf{b} - \text{diag}(\mathbf{X}\boldsymbol{\Phi}\mathbf{X}^\top)\|_2^2, \quad (7)$$

the first type is given by

$$\begin{aligned} \text{Type-A : } & \min_{\boldsymbol{\Omega}=[\mathbf{b}, \boldsymbol{\Phi}^\top]^\top \in \mathbb{R}^{(p+1) \times p}} \ell(\mathbf{b}, \boldsymbol{\Phi}) + \lambda_1 \|\boldsymbol{\Phi}\|_1 + \lambda_2 \sum_{j=1}^p \|[b_j, z(\boldsymbol{\phi}_j)]\|_q \\ & \text{s.t. } \boldsymbol{\Phi} = \boldsymbol{\Phi}^\top \text{ (for SH only),} \end{aligned} \quad (8)$$

where λ_1, λ_2 are regularization parameters, $1 < q \leq +\infty$ and $z(\mathbf{x})$ is a function satisfying the property that $z(\mathbf{x}) = \mathbf{0}$ implies $\mathbf{x} = \mathbf{0}$ for any vector $\mathbf{x} \in \mathbb{R}^p$. For instance, z can take the ℓ_r -norm function ($r > 0$)

$$z(\mathbf{x}) = \|\mathbf{x}\|_r, \quad (9)$$

or simply the identity function

$$z(\mathbf{x}) = \mathbf{x}^\top. \quad (10)$$

The first ℓ_1 penalty imposes elementwise sparsity on Φ and the second group- ℓ_1 penalty enforces column sparsity in Ω . We argue that with the two penalties and the constraint, (8) can be used for strong hierarchical variable selection. Indeed, the sparsity of \mathbf{b} comes from the second group-penalty alone, i.e., $b_j = 0$ implies $\|[b_j, z(\phi_j)]\|_q = 0$ (with probability 1) or $z(\phi_j) = 0$. By the property of the z function, $\phi_j = \mathbf{0}$, and thus $\phi_{jk} = 0$. The symmetry condition indicates further that $\phi_{kj} = 0$. Hence $(\phi_{jk} + \phi_{kj})/2$, the coefficient for $\mathbf{x}_j \odot \mathbf{x}_k$, is zero. Consequently, whenever $b_j = 0$, $\mathbf{x}_j \odot \mathbf{x}_k$ will be removed from the model and SH is automatically obeyed. We can describe the reasoning as follows

$$\begin{aligned} b_j = 0 &\Rightarrow \|[b_j, z(\phi_j)]\|_q = 0 \Rightarrow z(\phi_j) = 0 \Rightarrow \phi_j = \mathbf{0} \Rightarrow \phi_{jk} = 0 \\ &\Rightarrow \phi_{kj} = 0 \Rightarrow \frac{\phi_{jk} + \phi_{kj}}{2} = 0. \end{aligned} \quad (11)$$

Without the symmetry constraint, we can only complete the argument in the first line of (11), and so SH dose not hold. But interestingly, WH is guaranteed, because from $b_j = b_k = 0$, we have $(\phi_{jk} + \phi_{kj})/2 = 0$. Therefore, WH gives a relatively simpler problem.

As pointed out by a reviewer, when the model contains an intercept, centering the response and the p raw predictors does not make it vanish due to the presence of nonlinear terms, and so $\mathbf{1}b_0 + \mathbf{X}\mathbf{b} + \text{diag}(\mathbf{X}\Phi\mathbf{X}^\top)$ should be used to approximate \mathbf{y} . In the SH scenario, if at least one x -predictor is relevant, substituting $\mathbf{X}' = [\mathbf{1}, \mathbf{X}]$ for \mathbf{X} in (7) suffices.

We focus on convex forms of GRESH in this work. But surely the ℓ_1 penalty and the group- ℓ_1 penalty in (8) can be replaced by their nonconvex alternatives; see, e.g., She (2012).

GRESH is related to some methods in the literature. HL makes a special case of (8) because one of its formulations corresponds to $q = \infty$ and $r = 1$,

with a single regularization parameter being used. Another instance is given by $q = 2$ and $z(\mathbf{x}) = \mathbf{x}^\top$:

$$\min_{\boldsymbol{\Omega}=[\mathbf{b}, \boldsymbol{\Phi}^\top]^\top} \ell(\mathbf{b}, \boldsymbol{\Phi}) + \lambda_1 \|\boldsymbol{\Phi}\|_1 + \lambda_2 \|\boldsymbol{\Omega}^\top\|_{2,1} \quad \text{s.t. } \boldsymbol{\Phi} = \boldsymbol{\Phi}^\top. \quad (12)$$

Bien et al. (2013) incorrectly described VANISH (Radchenko and James, 2010) in this form, without the symmetry condition. We will focus on (12) in the theoretical and computational studies of Type-A GRESH.

As a matter of fact, Radchenko and James (2010) defined VANISH in a different way, which motivates another type of GRESH

$$\begin{aligned} \textbf{Type-B: } \min_{\boldsymbol{\Omega}=[\mathbf{b}, \boldsymbol{\Phi}^\top]^\top} & \ell(\mathbf{b}, \boldsymbol{\Phi}) + \lambda_1 \sum_{1 \leq j, k \leq p} \|\phi_{jk} \mathbf{x}_j \odot \mathbf{x}_k\|_2 + \\ & \lambda_2 \sum_{j=1}^p \|[b_j \mathbf{x}_j, z(\phi_{j1} \mathbf{x}_j \odot \mathbf{x}_1, \dots, \phi_{jp} \mathbf{x}_j \odot \mathbf{x}_p)]\|_q, \quad \text{s.t. } \boldsymbol{\Phi} = \boldsymbol{\Phi}^\top \text{ (SH only)}, \end{aligned} \quad (13)$$

where λ_1 , λ_2 , q , and z are defined as in (8). Similarly, we can argue that (13) keeps hierarchy. When $q = 2$ and z takes the form of (10), the penalty part in (13) become

$$\lambda_1 \sum_{j,k} \|\phi_{jk} \mathbf{x}_j \odot \mathbf{x}_k\|_2 + \lambda_2 \sum_j (\|b_j \mathbf{x}_j\|_2^2 + \sum_k \|\phi_{jk} \mathbf{x}_j \odot \mathbf{x}_k\|_2^2)^{\frac{1}{2}}, \quad (14)$$

as considered by Radchenko and James (2010). VANISH constructs main effects and interactions from two small sets of orthonormal basis functions in a functional regression setting. We do not pose such a restriction on the design matrix, and p can be arbitrarily large.

The key difference between the two types of GRESH is that the penalties are imposed on the coefficients in (8), but on the terms in (13). A common practice before calling a shrinkage method is normalizing/standardizing all predictors, so that it is more reasonable to use a common regularization parameter in penalizing different coefficients. In this way, (8) builds a model on the normalized predictors and their interactions, while Type-B amounts to forming the overall design $\tilde{\mathbf{X}}$ first and then performing the standardization. They are not equivalent because in general, $(\mathbf{x}_j/\|\mathbf{x}_j\|_2) \odot (\mathbf{x}_k/\|\mathbf{x}_k\|_2) \neq (\mathbf{x}_j \odot \mathbf{x}_k)/\|\mathbf{x}_j \odot \mathbf{x}_k\|_2$. Then, which type of GRESH is preferable? An answer will be given in Section 5.

GRESH offers some general schemes for hierarchical variable selection. But it is no ordinary lasso or group lasso, since Φ appears in both penalties as well as the symmetry constraint. The main goal of this paper is to tackle some computational and theoretical challenges arising from the overlapping regularization terms in high dimensions. In computation, we would like to develop fast and scalable algorithms (cf. Section 4); in theory, how to treat the penalties and the constraint *jointly* to derive a sharp error bound for GRESH is intriguing and challenging (cf. Sections 5, 6).

4 Computation

It is perhaps natural to think of using the alternating direction method of multipliers (ADMM, cf. Boyd et al. (2011)) to deal with the computational challenges. ADMM recently gains its popularity among statisticians. In fact, Bien et al. (2016) designed an algorithm of HL based on ADMM, where one of the main ingredients is the augmented Lagrangian

$$\begin{aligned} \min_{\mathbf{b}^{\pm} \in \mathbb{R}^p, \Phi \in \mathbb{R}^p} \quad & \ell(\mathbf{b}^+ - \mathbf{b}^-, \Phi) + \lambda \mathbf{1}^\top (\mathbf{b}^+ + \mathbf{b}^-) + \lambda \|\Phi\|_1 + \langle \Phi - \Psi, \mathbf{L} \rangle + \frac{\rho}{2} \|\Phi - \Psi\|_F^2, \\ \text{s.t.} \quad & \Psi = \Psi^\top, \|\phi_j\|_1 \leq b_j^+ + b_j^-, b_j^+ \geq 0, b_j^- \geq 0, 1 \leq j \leq p. \end{aligned} \quad (15)$$

Here, \mathbf{L} is a Lagrange multiplier matrix, and $\rho > 0$ is a given constant, sometimes referred as the penalty parameter. Although ADMM enjoys some nice convergence properties in theory, practically only when ρ is large enough can we obtain a solution with good statistical accuracy. But often the larger the value of ρ is, the slower the (primal) convergence is. For example, in the R package HierNet (version 1.6) for computing HL, $\rho = n$ is recommended, but for $p = 1000$, the algorithm may take several days to compute a single solution path. There are some empirical schemes on how to vary ρ during the iteration, but they are ad hoc and do not always behave well.

In this section, we consider a slightly more general optimization problem which includes both types of GRESH as particular instances

$$\begin{aligned} \min_{\Omega = [\mathbf{b}, \Phi^\top]^\top} \quad & \frac{1}{2} \|\mathbf{y} - \mathbf{X}\mathbf{b} - \text{diag}(\mathbf{Z}\Phi\mathbf{Z}^\top)\|_2^2 + \|\boldsymbol{\lambda}_b \odot \mathbf{b}\|_1 \\ & + \|\boldsymbol{\Lambda}_\Phi \odot \Phi\|_1 + \|\boldsymbol{\Lambda}_\Omega^\top \odot \Omega^\top\|_{2,1} \quad \text{s.t.} \quad \Phi = \Phi^\top, \end{aligned} \quad (16)$$

where $\mathbf{X}, \mathbf{Z} \in \mathbb{R}^{n \times p}$, and λ_b , Λ_Φ and Λ_Ω are non-negative regularization vector and matrices. Let $\bar{\mathbf{Z}} = [\mathbf{z}_1 \odot \mathbf{z}_1, \dots, \mathbf{z}_1 \odot \mathbf{z}_p, \dots, \mathbf{z}_p \odot \mathbf{z}_1, \dots, \mathbf{z}_p \odot \mathbf{z}_p]$, and $\check{\mathbf{Z}} = [\mathbf{x}_1, \mathbf{z}_1 \odot \mathbf{z}_1, \dots, \mathbf{z}_1 \odot \mathbf{z}_p, \dots, \mathbf{x}_p, \mathbf{z}_p \odot \mathbf{z}_1, \dots, \mathbf{z}_p \odot \mathbf{z}_p]$. We assume $\Lambda_\Omega = \mathbf{1}\lambda_\Omega^\top$ for some $\lambda_\Omega \in \mathbb{R}^p$ in developing the algorithm. (But since our algorithm applies to a general $\bar{\mathbf{Z}} \in \mathbb{R}^{n \times p^2}$ that is not necessarily symmetric, this is without loss of generality.) In (16), the ℓ_1 -type penalties are imposed on overlapping groups of variables. It is worth noting that the symmetry constraint considerably complicates the grouping structure. Without it, the variable groups can be shown to follow a tree structure, for which efficient algorithms can be developed on Jenatton et al. (2011) or Simon et al. (2013).

Our algorithm follows a different track than ADMM. The details are presented in Algorithm 1. Step 1 updates Ξ and results from a linearization-based surrogate function. Step 2 carries out a Dykstra-like splitting—see, e.g., Bauschke and Combettes (2008), by use of two proximity operators, $\vec{\Theta}_S$ and Θ_S . Concretely, for any real number a , $\Theta_S(a; \lambda)$ is given by $\text{sgn}(a)(|a| - \lambda)_+$ with $\text{sgn}(\cdot)$ representing the sign function. For any vector \mathbf{a} , $\Theta_S(\mathbf{a}; \lambda)$ is defined componentwise and the multivariate version $\vec{\Theta}_S(\mathbf{a}; \lambda)$ is given by $\mathbf{a}\Theta_S(\|\mathbf{a}\|_2; \lambda)/\|\mathbf{a}\|_2$ if $\mathbf{a} \neq \mathbf{0}$ and $\mathbf{0}$ otherwise.

The GRESH algorithm is easy to implement and involves no complicated matrix operations such as matrix inversion. Moreover, it does not contain ad-hoc algorithmic parameters like ρ in ADMM, and need no line search. Theorem 1 provides a universal choice for τ and guarantees the global optimality of $\hat{\Omega}$. In particular, strict iterate convergence, in addition to function-value convergence, can be established, which is considerably stronger than an “every accumulation point” type conclusion in many numerical studies. For clarity, we assume that the inner iteration runs till convergence, but this is unnecessary; see Remark 2 below.

Theorem 1. *Suppose $\lambda_b \geq \mathbf{0}$, $\Lambda_\Phi \geq \mathbf{0}$, $\lambda_\Omega > \mathbf{0}$. For any $\tau > \|\check{\mathbf{Z}}\|_2$ and any starting point $\Omega^{(0)}$, the sequence of iterates $\{\Omega^{(i)}\}$ converges to a globally optimal solution of (16).*

Remark 1. The conclusion in the theorem holds for “ $\ell_1 + \ell_2$ ” type penalties as well (Zou and Hastie, 2005; Owen, 2007). For the associated proximity operators, see He et al. (2013). In hierarchical variable selection, adding an ℓ_2 -type shrinkage is particularly helpful to compensate for model collinearity.

Remark 2. Neither the convergence of iterates nor the optimality guarantee requires the full convergence of the inner loop; see the Appendices A.1 for

Algorithm 1 The GRESH algorithm for solving the general problem (16)

Inputs:

Data: \mathbf{X} , \mathbf{Z} , \mathbf{y} . Regularization parameters: $\lambda_b, \Lambda_\Phi, \lambda_\Omega$.

Initialization:

$i \leftarrow 0$, τ large enough (say $\tau = \|\check{\mathbf{Z}}\|_2$);

$\lambda_b \leftarrow \lambda_b/\tau^2$, $\Lambda_\Phi \leftarrow (\Lambda_\Phi + \Lambda_\Phi^\top)/(2\tau^2)$, $\lambda_\Omega \leftarrow \lambda_\Omega/\tau^2$.

repeat

1. $\Xi_\Phi \leftarrow \Phi^{(i)} + \mathbf{Z}^\top \text{diag}\{\mathbf{y} - \mathbf{X}\mathbf{b}^{(i)} - \bar{\mathbf{Z}} \text{vec}(\Phi^{(i)})\} \mathbf{Z}/\tau^2$,
 $\Xi_b \leftarrow \mathbf{b}^{(i)} + \mathbf{X}^\top (\mathbf{y} - \mathbf{X}\mathbf{b}^{(i)} - \bar{\mathbf{Z}} \text{vec}(\Phi^{(i)}))/\tau^2$, $\Xi \leftarrow [\Xi_b, \Xi_\Phi]^\top$;
2. $\mathbf{P} \leftarrow \mathbf{0}$, $\mathbf{Q} \leftarrow \mathbf{0}$;

repeat

- (i) $\Omega[:, k] \leftarrow \vec{\Theta}_S(\Xi[:, k] + \mathbf{P}[:, k]; \lambda_\Omega[k]), \forall k : 1 \leq k \leq p$;
- (ii) $\mathbf{P} \leftarrow \mathbf{P} + \Xi - \Omega$;
- (iii) $\Xi[1, :] \leftarrow \Theta_S(\Omega[1, :] + \mathbf{Q}[1, :]; \lambda_b^\top)$;
- (iv) $\Omega[2:\text{end}, :] \leftarrow (\Omega[2:\text{end}, :] + \Omega[2:\text{end}, :]^\top)/2$;
- (v) $\Xi[2:\text{end}, :] \leftarrow \Theta_S(\Omega[2:\text{end}, :] + \mathbf{Q}[2:\text{end}, :]; \Lambda_\Phi)$;
- (vi) $\mathbf{Q} \leftarrow \mathbf{Q} + \Omega - \Xi$;

until convergence

3. $\Omega^{(i+1)} \leftarrow \Xi$;

4. $i \leftarrow i + 1$;

until convergence

Output $\hat{\Omega}$

more detail. Various stopping criteria can be employed, e.g., Schmidt et al. (2011). In our experience, running (i)–(vi) for a few steps (say 10) usually suffices.

Remark 3. Algorithm 1 and Theorem 1 can be extended beyond quadratic loss functions. When ℓ takes the binomial deviance in classification problems, the first step of Algorithm 1 becomes $\Xi_b \leftarrow \mathbf{b}^{(i)} + \mathbf{X}^\top (\mathbf{y} - \pi(\mathbf{X}\mathbf{b}^{(i)} + \bar{\mathbf{Z}} \text{vec}(\Phi^{(i)})))/\tau^2$, $\Xi_\Phi \leftarrow \Phi^{(i)} + \mathbf{Z}^\top \text{diag}\{\mathbf{y} - \pi(\mathbf{X}\mathbf{b}^{(i)} + \bar{\mathbf{Z}} \text{vec}(\Phi^{(i)}))\} \mathbf{Z}/\tau^2$, $\Xi \leftarrow [\Xi_b, \Xi_\Phi]^\top$, where $\pi(t) = 1/(1 + \exp(-t))$ and extends componentwise to vectors. We can show that choosing $\tau > \|\check{\mathbf{Z}}\|_2/2$ guarantees the convergence of the algorithm.

Remark 4. We recommend applying Nesterov’s first acceleration in implementations (Nesterov, 2007). In more detail, it uses a momentum update

of Ξ in Step 1: If $i = 0$, $\Xi_b \leftarrow \mathbf{b}^{(i)} + \mathbf{X}^\top(\mathbf{y} - \mathbf{X}\mathbf{b}^{(i)} - \bar{\mathbf{Z}} \text{vec}(\Phi^{(i)}))/\tau^2$, $\Xi_\Phi \leftarrow \Phi^{(i)} + \mathbf{Z}^\top \text{diag}\{\mathbf{y} - \mathbf{X}\mathbf{b}^{(i)} - \bar{\mathbf{Z}} \text{vec}(\Phi^{(i)})\}\mathbf{Z}/\tau^2$; if $i > 0$, $\Xi_b \leftarrow (1 - \omega_i)\Xi_b + \omega_i(\mathbf{b}^{(i)} + \mathbf{X}^\top(\mathbf{y} - \mathbf{X}\mathbf{b}^{(i)} - \bar{\mathbf{Z}} \text{vec}(\Phi^{(i)}))/\tau^2)$, $\Xi_\Phi \leftarrow (1 - \omega_i)\Xi_\Phi + \omega_i(\Phi^{(i)} + \mathbf{Z}^\top \text{diag}\{\mathbf{y} - \mathbf{X}\mathbf{b}^{(i)} - \bar{\mathbf{Z}} \text{vec}(\Phi^{(i)})\}\mathbf{Z}/\tau^2)$, where $\omega_i = (2i+3)/(i+3)$. Empirically, the number of iterations can be reduced by about 40% in comparison to the non-relaxed form.

5 Non-asymptotic analysis

In this section, given any $\Delta \in \mathbb{R}^{(p+1) \times p}$ and $\mathcal{J}_G \subset [p]$, we use $\Delta_{\mathcal{J}_G}$ to denote the submatrix $\Delta[:, \mathcal{J}_G]$. Given any $\mathcal{J}_e \subset [p^2]$, $\|\text{vec}(\Delta_\Phi)_{\mathcal{J}_e}\|_1$ and $\|\text{vec}(\Delta_\Phi)_{\mathcal{J}_e}\|_2$ are abbreviated as $\|(\Delta_\Phi)_{\mathcal{J}_e}\|_1$ and $\|(\Delta_\Phi)_{\mathcal{J}_e}\|_2$, respectively, when there is no ambiguity.

In this multi-regularization setting, standard treatments of the stochastic term do not give sharp error rates. In particular, applying $\langle \varepsilon, \bar{\mathbf{X}} \text{vec}(\Delta_\Phi) \rangle \leq \|\bar{\mathbf{X}}^\top \varepsilon\|_\infty \|\Delta_\Phi\|_1$ and $\langle \varepsilon, \check{\mathbf{X}} \text{vec}(\Delta) \rangle \leq \|\check{\mathbf{X}}^\top \varepsilon\|_{2,\infty} \|\Delta^\top\|_{2,1}$ as commonly used in the literature (Bickel et al., 2009; Lounici et al., 2011; Negahban et al., 2012; van de Geer, 2014) would yield a prediction error bound of the order $\sigma^2(J_e \log p + J_G p)$, which is, ironically, much worse than the error rate of LASSO or Group LASSO (G-LASSO). Our analysis relies on two interrelated inequalities derived from the statistical and computational properties of GRESH estimators. See the Appendices for more technical detail.

First, let's consider the **Type-A** problem (12), with λ_1, λ_2 redefined:

$$\min_{\Omega=[\hat{\mathbf{b}}, \hat{\Phi}^\top]^\top} \frac{1}{2} \|\mathbf{y} - \mathbf{X}\hat{\mathbf{b}} - \bar{\mathbf{X}} \text{vec}(\Phi)\|_2^2 + \lambda_1 \|\check{\mathbf{X}}\|_2 \|\Phi\|_1 + \lambda_2 \|\check{\mathbf{X}}\|_2 \|\Omega^\top\|_{2,1} \text{ s.t. } \Phi = \Phi^\top. \quad (17)$$

Let $\hat{\Omega} = [\hat{\mathbf{b}}, \hat{\Phi}^\top]^\top$ be any global minimizer of (17). We are interested in its prediction accuracy measured by $M(\hat{\mathbf{b}} - \mathbf{b}^*, \hat{\Phi} - \Phi^*)$, where

$$M(\mathbf{b}, \Phi) = \|\mathbf{X}\mathbf{b} + \bar{\mathbf{X}} \text{vec}(\Phi)\|_2^2. \quad (18)$$

The predictive learning perspective is always legitimate in evaluating the quality of the estimator regardless of the signal-to-noise ratio. To guarantee small predictor errors when using a convex method, the design matrix must satisfy certain incoherence conditions, one of the most popular being the restricted eigenvalue (RE) assumption (Bickel et al., 2009; Lounici et al., 2011). In the following, we give an extension of RE in the hierarchy setting,

with the restricted cone defined with both ℓ_1 and group- ℓ_1 penalties. A less intuitive but technically much less demanding condition is used in the proof.

ASSUMPTION $\mathcal{A}(\mathcal{J}_e, \mathcal{J}_G, \vartheta, \kappa)$. Given $\mathcal{J}_e \subset [p^2]$, $\mathcal{J}_G \subset [p]$, $\kappa \geq 0$ and a constant $\vartheta \geq 0$, for any $\Delta = [\Delta_b, \Delta_\Phi]^\top \in \mathbb{R}^{(p+1) \times p}$ satisfying $\Delta_\Phi = \Delta_\Phi^\top$ and $\|(\Delta_\Phi)_{\mathcal{J}_e^c}\|_1 + \|(\Delta_{\mathcal{J}_G^c})^\top\|_{2,1} \leq (1 + \vartheta)(\|(\Delta_\Phi)_{\mathcal{J}_e}\|_1 + \|(\Delta_{\mathcal{J}_G})^\top\|_{2,1})$, the following inequality holds

$$\kappa \|\check{\mathbf{X}}\|_2^2 (\|(\Delta_\Phi)_{\mathcal{J}_e}\|_2^2 + \|\Delta_{\mathcal{J}_G}\|_F^2) \leq \|\mathbf{X} \Delta_b + \bar{\mathbf{X}} \text{vec}(\Delta_\Phi)\|_2^2. \quad (19)$$

The rate choices of the regularization parameters play a major role in prediction. We choose λ_1 and λ_2 in (17) according to

$$\lambda_1 = A_1 \sigma \sqrt{\log(ep)}, \quad \lambda_2 = A_2 \sigma \sqrt{\log(ep)}, \quad (20)$$

where A_1, A_2 are large constants. (20) is quite different from the typical choice in group- ℓ_1 penalization; see Remark 2 for more detail.

The following theorem states a non-asymptotic oracle inequality as well as a model cardinality bound for GRESH estimators. For convenience, we use abbreviated symbols $\hat{J}_G = J_G(\hat{\mathbf{b}}, \hat{\Phi})$ and $\hat{J}_e = J_e(\hat{\Phi})$ for the estimate, and $J_G = J_G(\mathbf{b}, \Phi)$ and $J_e = J_e(\Phi)$ for the reference signal.

Theorem 2. Assume $\epsilon \sim N(\mathbf{0}, \sigma^2 \mathbf{I})$. Let $\hat{\Omega} = [\hat{\mathbf{b}}, \hat{\Phi}^\top]^\top$ be a global minimizer of (17). Then under (20), for any sufficiently large constants A_1, A_2 , the following oracle inequality holds for any $(\mathbf{b}, \Phi) \in \mathbb{R}^p \times \mathbb{R}^{p \times p}$

$$\mathbb{E}[M(\hat{\mathbf{b}} - \mathbf{b}^*, \hat{\Phi} - \Phi^*)] \lesssim M(\mathbf{b} - \mathbf{b}^*, \Phi - \Phi^*) + (1 \vee \frac{1}{\kappa}) \sigma^2 (J_e + J_G) \log p + \sigma^2, \quad (21)$$

provided that $(\mathbf{X}, \bar{\mathbf{X}}, \check{\mathbf{X}})$ satisfies $\mathcal{A}(\mathcal{J}_e, \mathcal{J}_G, \vartheta, \kappa)$ for some $\kappa > 0$ and some constant $\vartheta \geq 0$. Furthermore, under the same regularity condition, the overall sparsity of the obtained model is controlled by

$$\mathbb{E}[\hat{J}_e] + \mathbb{E}[\hat{J}_G] \lesssim \{M(\mathbf{b}^* - \mathbf{b}, \Phi^* - \Phi) + \sigma^2\} / \{\sigma^2 \log(ep)\} + J_e + J_G. \quad (22)$$

Remark 1. Letting $\mathbf{b} = \mathbf{b}^*$ and $\Phi = \Phi^*$ in (21), we obtain an error bound no larger than $\sigma^2 (J_e^* + J_G^*) \log p$ (omitting constant factors). This indicates that GRESH not only guarantees SH, but can give an error rate as low as that of LASSO. The existence of the bias term $M(\mathbf{b} - \mathbf{b}^*, \Phi - \Phi^*)$ makes our results applicable to approximately sparse signals, which is of practical

significance. The theorem does not require the spectral norms of the design matrices \mathbf{X} , $\bar{\mathbf{X}}$ and $\check{\mathbf{X}}$ to be bounded above by $O(\sqrt{n})$ as assumed in, for example, Zhang and Huang (2008) and Bickel et al. (2009). In addition, the true signal $\boldsymbol{\Omega}^*$ and the reference signal $\boldsymbol{\Omega}$ in the theorem need not obey SH.

Remark 2. It is widely acknowledged that the penalty parameter for a grouped ℓ_1 penalty should be adjusted by the group size (Yuan and Lin, 2006). In fact, λ_2 would be of order $\sigma\sqrt{p + \log p}$ from Lounici et al. (2011) and Wei and Huang (2010), in light of the fact that there are p groups of size $(p + 1)$ in $\|\boldsymbol{\Omega}\|_{2,1}$. Perhaps surprisingly, this parameter choice becomes suboptimal in hierarchical variable selection. In fact, due to the presence of multiple penalties, we show in the proof that (20) suffices to suppress the noise, which in turn leads to a reduced error rate. Such a novel finding is owing to the careful treatment of the stochastic term, which is generally applicable to overlap group lasso (Jenatton et al., 2011). The conclusion that λ_1 and λ_2 are of essentially the same rate also facilitates parameter tuning, since one just needs to search along a one-dimensional grid.

Remark 3. Theorem 2 can be extended to sub-Gaussian $\text{vec}(\boldsymbol{\varepsilon})$ with mean 0 and its ψ_2 -norm bounded by σ , which covers more noise distributions. High probability form results of the prediction error can be obtained as well: (21) and (22), without the expectation and the additive σ^2 term, hold with probability at least $1 - Cp^{-c\min\{A_1^2, A_2^2\}}$ for some universal constants C and c , and so $\hat{J}_e + \hat{J}_G \lesssim J_e^* + J_G^*$ with high probability. Moreover, in the Appendices B, we show how to adapt our proof to deliver a coordinatewise error bound which can be used for recovering the sparsity pattern of the true signal.

Remark 4. For the WH version of (17) (without the symmetry condition), following the lines of the proof of Theorem 2, we can show its error rate is of the order $\sigma^2\{J_e^w(\boldsymbol{\Phi}) + J_G^w(\boldsymbol{\Omega})\} \log p$, where $J_G^w(\boldsymbol{\Omega}) = J_G(\boldsymbol{\Omega}')$, $J_e^w(\boldsymbol{\Phi}) = J_e(\boldsymbol{\Phi}')$, with $\boldsymbol{\Omega}' = [\mathbf{b}, \boldsymbol{\Phi}'^\top]^\top$ and $\phi'_{kj} = \phi_{kj} + \phi_{jk}$ for $k \geq j$ and 0 otherwise. The associated regularity condition uses $\mathcal{J}_e^w, \mathcal{J}_G^w$, in place of $\mathcal{J}_e, \mathcal{J}_G$, respectively, and does not require $\boldsymbol{\Delta}_\Phi$ to be symmetric. Details are not reported in the paper.

Similarly, we can derive an oracle inequality for GRESH estimators of Type-B. Let \mathbf{X}^s be the column-scaled \mathbf{X} such that the ℓ_2 -norm of each of its columns equals 1. $\bar{\mathbf{X}}^s$ and $\check{\mathbf{X}}^s$ are similarly defined. The corresponding coefficients, denoted by \mathbf{b}^s and $\boldsymbol{\Phi}^s$, satisfy $\mathbf{X}\mathbf{b} = \mathbf{X}^s\mathbf{b}^s$, $\bar{\mathbf{X}}\text{vec}(\boldsymbol{\Phi}) = \bar{\mathbf{X}}^s\text{vec}(\boldsymbol{\Phi}^s)$. Let $(\hat{\mathbf{b}}^s, \hat{\boldsymbol{\Phi}}^s)$ be a global minimizer of the scaled **Type-B**

problem: $\min_{\Omega^s = [\mathbf{b}^s, \Phi^s]^\top} \frac{1}{2} \|\mathbf{y} - \mathbf{X}^s \mathbf{b}^s - \bar{\mathbf{X}}^s \text{vec}(\Phi^s)\|_2^2 + \lambda_1 \|\check{\mathbf{X}}^s\|_2 \|\Phi^s\|_1 + \lambda_2 \|\check{\mathbf{X}}^s\|_2 \|(\Omega^s)^\top\|_{2,1}$ s.t. $\Phi^s = (\Phi^s)^\top$. As aforementioned, the problem can not be reduced to (17) because $\bar{\mathbf{X}}^s[:, jk]$ does not equal $\mathbf{X}^s[:, j] \odot \mathbf{X}^s[:, k]$ in general.

ASSUMPTION $\mathcal{A}'(\mathcal{J}_e, \mathcal{J}_G, \vartheta, \kappa')$. Given $\mathcal{J}_e \subset [p^2]$, $\mathcal{J}_G \subset [p]$, and positive constants κ' and ϑ , for any $\Delta = [\Delta_b, \Delta_\Phi^\top]^\top$ satisfying $\Delta_\Phi = \Delta_\Phi^\top$ and $\|(\Delta_\Phi)_{\mathcal{J}_e^c}\|_1 + \|(\Delta_{\mathcal{J}_G^c})^\top\|_{2,1} \leq (1 + \vartheta)(\|(\Delta_\Phi)_{\mathcal{J}_e}\|_1 + \|(\Delta_{\mathcal{J}_G})^\top\|_{2,1})$, the following inequality holds

$$\kappa' \|\check{\mathbf{X}}^s\|_2^2 (\|(\Delta_\Phi)_{\mathcal{J}_e}\|_2^2 + \|\Delta_{\mathcal{J}_G}\|_F^2) \leq \|\mathbf{X}^s \Delta_b + \bar{\mathbf{X}}^s \text{vec}(\Delta_\Phi)\|_2^2.$$

Theorem 2' *Under the same conditions as in Theorem 2 and with $\mathcal{A}'(\mathcal{J}_e, \mathcal{J}_G, \vartheta, \delta'_{\mathcal{J}_e, \mathcal{J}_G})$ in place of $\mathcal{A}(\mathcal{J}_e, \mathcal{J}_G, \vartheta, \kappa)$, (21) and (22) hold.*

The error bounds of the two types of GRESH are of the same order, but their regularity conditions place different requirements on the design. We performed extensive simulation studies to compare \mathcal{A} and \mathcal{A}' , and found that for the same ϑ , $\kappa < \kappa'$ usually holds, which suggests the penalization on the basis of terms seems more appropriate than that on the coefficients. Therefore, we recommend Type-B regularization for hierarchical variable selection.

6 Minimax lower bound and error rate comparison

In this section, we show that in a minimax sense, the error rate we obtained in Theorem 2 is minimax optimal up to some logarithm factors. Consider two signal classes having hierarchy and joint sparsity:

$$\text{SH}(J_G, J_e) = \{\Omega = [\mathbf{b}, \Phi^\top]^\top : \Omega \text{ obeys SH}, \Phi = \Phi^\top, J_G(\Omega) \leq J_G, J_e(\Phi) \leq J_e\}, \quad (23)$$

$$\text{WH}(J_G, J_e) = \{\Omega = [\mathbf{b}, \Phi^\top]^\top : \Omega \text{ obeys WH}, J_G^w(\Omega) \leq J_G, J_e^w(\Phi) \leq J_e\}, \quad (24)$$

where $1 \leq J_G \leq p$, $1 \leq J_e \leq pJ_G$. Recall the definitions of J_G^w and J_e^w in Remark 4 following Theorem 2. Let $\ell(\cdot)$ be a nondecreasing loss function

with $\ell(0) = 0$, $\ell \neq 0$. Under some regularity assumptions, we study the minimax lower bounds for strong and weak hierarchical variable selection.

ASSUMPTION $\mathcal{B}^S(J_G, J_e)$. For any $\mathbf{\Omega} = [\mathbf{b}, \mathbf{\Phi}^\top]^\top \in \mathbb{R}^{(p+1) \times p}$ satisfying that $\mathbf{\Phi}$ is symmetric, $J_e(\mathbf{\Phi}) \leq J_e$ and $J_G(\mathbf{\Omega}) \leq J_G$, $\underline{\kappa} \|\check{\mathbf{X}}\|_2^2 \|\mathbf{\Omega}\|_F^2 \leq \|\check{\mathbf{X}} \text{vec}(\mathbf{\Omega})\|_2^2 \leq \bar{\kappa} \|\check{\mathbf{X}}\|_2^2 \|\mathbf{\Omega}\|_F^2$ holds for some positive constants $\underline{\kappa}$ and $\bar{\kappa}$.

ASSUMPTION $\mathcal{B}^W(J_G, J_e)$. For any $\mathbf{\Omega} = [\mathbf{b}, \mathbf{\Phi}^\top]^\top \in \mathbb{R}^{(p+1) \times p}$ satisfying that $J_e^w(\mathbf{\Phi}) \leq J_e$ and $J_G^w(\mathbf{\Omega}) \leq J_G$, $\underline{\kappa} \|\check{\mathbf{X}}\|_2^2 \|\mathbf{\Omega}\|_F^2 \leq \|\check{\mathbf{X}} \text{vec}(\mathbf{\Omega})\|_2^2 \leq \bar{\kappa} \|\check{\mathbf{X}}\|_2^2 \|\mathbf{\Omega}\|_F^2$ holds for some positive constants $\underline{\kappa}$ and $\bar{\kappa}$.

Theorem 3. (i) *Strong hierarchy.* Assume $\mathbf{y} = \mathbf{X}\mathbf{b}^* + \bar{\mathbf{X}} \text{vec}(\mathbf{\Phi}^*) + \boldsymbol{\varepsilon}$ with $\boldsymbol{\varepsilon} \sim N(\mathbf{0}, \sigma^2 \mathbf{I})$, $J_G \geq 2$, $p \geq 2$, $J_e \geq 1$, $n \geq 1$, $J_G \leq p/2$, $J_e \leq J_G^2/2$, and $\mathcal{B}^S(2J_G, 2J_e)$ is satisfied. Then there exist positive constants C, c (depending on $\ell(\cdot)$ only) such that

$$\inf_{\hat{\mathbf{\Omega}}} \sup_{\mathbf{\Omega}^* \in SH(J_G, J_e)} \mathbb{E}[\ell(M(\hat{\mathbf{b}} - \mathbf{b}^*, \hat{\mathbf{\Phi}} - \mathbf{\Phi}^*) / (CP_o(J_e, J_G)))] \geq c > 0, \quad (25)$$

where $\hat{\mathbf{\Omega}}$ denotes any estimator, and

$$P_o(J_e, J_G) = \sigma^2 \{J_e \log(eJ_G^2/J_e) + J_G \log(ep/J_G)\}. \quad (26)$$

(ii) *Weak hierarchy.* Let $J_G \geq 1$, $J_e \geq 1$, $n \geq 1$, $p \geq 2$, $J_G \leq p/2$, $J_e \leq J_G p/2$. Under the same model assumption and $\mathcal{B}^W(2J_G, 2J_e)$, (25) holds if $SH(J_G, J_e)$ is replaced by $WH(J_G, J_e)$ and P_o is replaced by

$$P'_o(J_e, J_G) = \sigma^2 \{J_e \log(eJ_G p/J_e) + J_G \log(ep/J_G)\}. \quad (27)$$

We give some examples of ℓ to illustrate the conclusion. For SH, using the indicator function $\ell(u) = 1_{u \geq 1}$, we know that for any estimator $(\hat{\mathbf{b}}, \hat{\mathbf{\Phi}})$,

$$M(\hat{\mathbf{b}} - \mathbf{b}^*, \hat{\mathbf{\Phi}} - \mathbf{\Phi}^*) \gtrsim \sigma^2 (J_e \log(eJ_G^2/J_e) + J_G \log(ep/J_G))$$

occurs with positive probability, under some mild conditions. For $\ell(u) = u$, Theorem 3 shows that the risk $\mathbb{E}[M(\hat{\mathbf{b}} - \mathbf{b}^*, \hat{\mathbf{\Phi}} - \mathbf{\Phi}^*)]$ is bounded from below by $P_o(J_e, J_G)$ up to some multiplicative constant. Because $J_G \leq J_e \leq J_G p$ and $J_G \leq p$, it is easy to see that the minimax rates are no larger than the error rate obtained in Theorem 2.

A comparison of some popular methods follows, where we can see the benefits of hierarchical variable selection. LASSO, in our context, solves $\min_{\mathbf{b}, \Phi: \Phi = \Phi^\top} \ell(\mathbf{b}, \Phi) + \lambda(\|\mathbf{b}\|_1 + \|\Phi\|_1)$. From Bickel et al. (2009), the estimator has a prediction error of the order $\sigma^2(J_e + J^{10} + J^{11}) \log p$. G-LASSO, with the optimization problem defined by $\min_{\mathbf{b}, \Phi: \Phi = \Phi^\top} \ell(\mathbf{b}, \Phi) + \lambda\|\Omega^\top\|_{2,1}$, automatically maintains SH and has an error rate of $\sigma^2 J_G p$ (Lounici et al., 2011). In general, there is no clear winner between the two. Let's turn to a particularly interesting case where $J_e^* \ll J_G^* p$, i.e., the existence of a main effect in the model does not indicate that all its associated interactions must be relevant. In this scenario, LASSO always outperforms G-LASSO, although it does not possess the SH property. GRESH achieves the same low error rate and guarantees hierarchy, because under SH, $J^{11} + J^{10} = J_G$.

The error rate proved in (21) does not always beat that of G-LASSO, because only large values of A are considered in Theorem 2. Yet, even in the worst case when $J_e^* \asymp J_G^* p$, GRESH is only a logarithmic factor worse. In practical data analysis, there will be no performance loss, because when $\lambda_1 = 0$, GRESH degenerates to G-LASSO.

Table 1: Error rate comparison between LASSO, G-LASSO, and GRESH, where σ^2 and other constant factors are omitted.

LASSO	$(J_e + J^{11} + J^{10}) \log p$
G-LASSO	$J_G p$
GRESH	$(J_G + J_e) \log p$
Minimax:	$J_G \log(ep/J_G) + J_e \log(eJ_G^2/J_e)$

7 Experiments

7.1 Simulations

In this part, we perform some simulation studies to compare the performance of HL and GRESH (of Type B, cf. (14)) in terms of prediction accuracy, selection consistency, and computational efficiency. We use a Toeplitz design to generate all main predictors, with the correlation between \mathbf{x}_i and \mathbf{x}_j given by $0.5^{|i-j|}$. The true coefficients \mathbf{b}^* and Φ^* (symmetric) are generated according to the following three setups.

Example 1. $n = 40$, $p = 100$ or 200 (and so $p + \binom{p}{2} = 5050$ or 20100). $\mathbf{b}^* = [3, 1.5, 0, 0, 2, 2, 0, \dots, 0]^\top$, $\Phi^* = \mathbf{0}$, $\sigma^2 = 1$. No interactions are relevant to the response variable. SH is satisfied.

Example 2. $n = 150$, $p = 50$ or 100 (and so $p + \binom{p}{2} = 1275$ or 5050). $\mathbf{b}^* = [3, 3, 3, 3, 3, 3, 3, 3, 0, \dots, 0]^\top$, $\Phi^*[1, 2] = \Phi^*[1, 3] = \Phi^*[4, 5] = \Phi^*[4, 6] = \Phi^*[7, 8] = \Phi^*[7, 9] = 3$, $\sigma^2 = 1$. The model involves both main and interaction effects and obeys SH.

Example 3. $n = 100$, $p = 50$ or 100 (and so $p + \binom{p}{2} = 1275$ or 5050). $\mathbf{b}^* = [1, 1, 1, 1, 1, 1, 1, 0, \dots, 0]^\top$, $\Phi^*[i, j] = 5$, $1 \leq i, j \leq 5$, $i \neq j$, $\Phi^*[4, 5] = \Phi^*[4, 6] = \Phi^*[4, 7] = 5$, $\sigma^2 = 1$. The true model does not have very strong main effects but satisfies SH.

All regularization parameters are tuned on a (separate) large validation dataset containing 10K observations. There is no need to perform a full two-dimensional grid search to find the optimal parameters in GRESH. Rather, motivated by Theorem 2, we set $\lambda_2 = c\lambda_1$, and chose $c = 0.5$ according to experience. Because of the convex nature of the problem, pathwise computation with warm starts is used. After variable selection, a ridge regression model is always refitted to be used for prediction. The official R package for HL is HierNet, implemented in C. We set `strong=TRUE` and post-calibrate HL by a restricted ridge refitting, which substantially enhances its accuracy. To make a fair comparison between HL and GRESH, we use the same error tolerance ($1e-5$) and the same number of grid values (20). All other algorithmic parameters in HierNet are set to their default values. Given each setup, we repeat the experiment for 50 times and evaluate the performance of each algorithm according to the measures defined below. The test error (Err) is the mean squared error between the true mean of \mathbf{y} and its estimate; for robustness and stability, we report the median test error from all runs. The joint detection (JD) rate is the fraction of $|\{(i, j) : \Omega_{ij}^* \neq 0\}| \subseteq |\{(i, j) : \hat{\Omega}_{ij} \neq 0\}|$ among all experiments. The missing (M) rate and the false alarm (FA) rate are the mean of $|\{(i, j) : \Omega_{ij}^* \neq 0, \hat{\Omega}_{ij} = 0\}| / |\{(i, j) : \Omega_{ij}^* \neq 0\}|$ and the mean $|\{(i, j) : \Omega_{ij}^* = 0, \hat{\Omega}_{ij} \neq 0\}| / |\{(i, j) : \Omega_{ij}^* = 0\}|$, respectively. The path computational cost is the average running time of an algorithm in seconds. All the experiments were run on a PC with 3.2GHz CPU, 32GB memory and 64-bit Windows 8.1. Table 2 and Table 3 summarize the statistical and computational results.

From Table 2, GRESH and HL behaved equally well in Example 1, the

Table 2: Statistical performance of HL and GRESH, measured in test error, joint detection rate, missing rate, and false alarm rate on simulation data. All numbers are multiplied by 100.

	Ex 1 ($p = 100$)				Ex 2 ($p = 50$)				Ex 3 ($p = 50$)			
	Err	JD	M	FA	Err	JD	M	FA	Err	JD	M	FA
HL	13.7	100	0.00	0.00	14.2	95	0.24	0.31	68.0	90	0.65	2.71
GRESH	13.6	100	0.00	0.00	11.6	100	0.00	0.12	26.2	100	0.00	0.23
	Ex 1 ($p = 200$)				Ex 2 ($p = 100$)				Ex 3 ($p = 100$)			
	Err	JD	M	FA	Err	JD	M	FA	Err	JD	M	FA
HL	13.8	100	0.00	0.00	17.3	90	0.48	0.11	92.2	15	24.03	1.44
GRESH	13.5	100	0.00	0.00	14.3	100	0.00	0.05	26.9	100	0.00	0.08

Table 3: Path computation costs of GRESH and HL when $p = 200$ and 1000. The computational times are in seconds unless otherwise specified.

	$p = 200$			$p = 1000$		
	Ex1	Ex2	Ex3	Ex1	Ex2	Ex3
HL	574	3057	2.9 hours	7.6 hours	—	—
GRESH	110	158	128	1066	2.5 hours	3194

model of which contains main effects only, but GRESH is faster. In Example 2 and Example 3, the two methods show more differences; see their test errors and joint identification rates, for example. We also noticed that GRESH often gave a more parsimonious model. When the main effects are weak as in Example 3, HL may miss some genuine interaction effects. Overall, GRESH showed comparable or better test errors. In fact, this is observed even when SH is not satisfied (results not shown in the table). We suspect that the performance differences between HL and GRESH largely result from the fact that HL compares $|b_j|$ with $\|\phi_j\|_1$, the ℓ_1 -norm of the overall ϕ_j , to realize SH, while (14) groups $b_j\mathbf{x}_j$, $\phi_{j1}\mathbf{x}_j \odot \mathbf{x}_1$, \dots , and $\phi_{jp}\mathbf{x}_j \odot \mathbf{x}_p$, on the term basis, to select main effects.

The computational times in Table 3 show the scalability of each algorithm as p varies. When $p = 1000$, there are 500500 variables in total, and so HL became computationally prohibitive, also evidenced by Lim and Hastie (2015). GRESH offered impressive computation gains in the experiment.

7.2 Comparison with ADMM

This part shows that directly applying ADMM does not give a scalable algorithm for solving the optimization problem (16) which has a large number of groups with large group size. The detailed algorithm design is given in the Appendices. We set $\rho = 1$ in ADMM and compared it to Algorithm 1. The results are reported in Table 4 and Table 5. The statistical performances of the two algorithms are close. This is reasonable because they solve the same optimization problem. However, ADMM is much slower. In the experiments, ADMM became infeasible when $p = 200$ or larger.

Table 4: Statistical performance of GRESH and ADMM, measured in test error, joint detection rate, missing rate, and false alarm rate.

	Ex 1 ($p = 50$)				Ex 2 ($p = 50$)				Ex 3 ($p = 50$)			
	Err	JD	M	FA	Err	JD	M	FA	Err	JD	M	FA
GRESH	11.6	100	0.00	0.02	11.6	100	0.00	0.12	26.2	100	0.00	0.23
ADMM	11.7	100	0.00	0.02	11.9	100	0.00	0.12	26.3	100	0.00	0.22
	Ex 1 ($p = 100$)				Ex 2 ($p = 100$)				Ex 3 ($p = 100$)			
	Err	JD	M	FA	Err	JD	M	FA	Err	JD	M	FA
GRESH	13.6	100	0.00	0.00	14.3	100	0.00	0.05	26.9	100	0.00	0.08
ADMM	16.8	100	0.00	0.01	14.5	100	0.00	0.06	27.1	100	0.00	0.09

Table 5: Path computation costs of GRESH and ADMM.

	$p = 50$			$p = 100$		
	Ex1	Ex2	Ex3	Ex1	Ex2	Ex3
GRESH	8	11	9	27	36	30
ADMM	66	70	60	1027	1571	1439

7.3 Real data example

We performed hierarchical variable selection on the California housing data (Pace and Barry, 1997). The dataset consists of 9 summary characteristics for 20640 neighborhoods in California. The response variable is the median

house value in each neighborhood. Following Hastie et al. (2009), we obtained eight household-related predictor variables: median income, housing median age, average number of rooms and bedrooms per household, population, average occupancy (population/households), latitude, and longitude, denoted by `MedInc`, `Age`, `AvgRms`, `AvgBdrms`, `Popu`, `AvgOccu`, `Lat`, and `Long`, respectively. Similar to Ravikumar et al. (2007) and Radchenko and James (2010), 50 nuisance features generated as standard Gaussian random variables were added, to make the problem more challenging. The full quadratic model on this enlarged dataset contains 3422 unknowns.

To prevent from getting over-optimistic error estimates, we used a hierarchical cross-validation procedure where an outer 10-fold cross-validation (CV) is for performance evaluation and the inner 10-fold CVs are for parameter tuning. We managed to run both HL and GRESH for hierarchical variable selection, with the estimates post-calibrated by a local ridge fitting as described in Section 7.1. It took us approximately one and half days to complete the CV experiment for HL, and about 1.6 hours for GRESH. The median and mean test errors of the models obtained by HL are 530.8 and 553.5, respectively, and the average number of selected variables is 31.2. GRESH gave 516.9 and 521.1 for the median and mean test errors, respectively, and selected 17.1 variables on average, about half of the model size of HL.

To help the reader get an intuition of the selection frequencies of all predictors, we display heat maps in Figure 1. The two heat maps in the top panel include all variables, and the bottom panel only shows the heat maps restricted to the original covariates and their interactions. According to the figure, both methods successfully removed most of the artificially added noisy features. On average, only 9.3 nuisance covariates exist in the models obtained by HL, and 5.4 in the GRESH models. The heat maps of GRESH are however neater. The HL selection results are less parsimonious and are perhaps more difficult to interpret. The nonlinear terms in GRESH include the interaction between `MedInc` and `Age`, in addition to the quadratic effects of `MedInc`, `Age`, and `AvgBdrms`. `Popu` and all its associated interaction terms never got selected by GRESH. The insignificance of `Popu` can be confirmed by more elaborate analysis based on gradient boosting (Hastie et al., 2009).

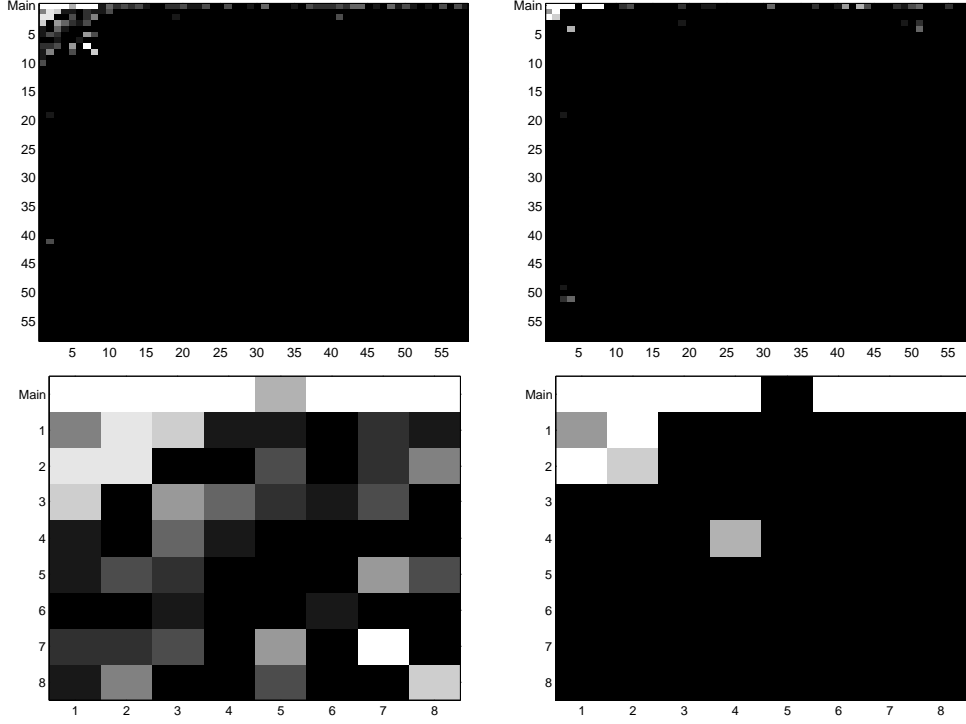


Figure 1: Top panel: heat maps of HL (left) and GRESH (right) on California housing data. Bottom panel: heat maps of HL (left) and GRESH (right) restricted to the original 8 variables and their interactions.

Appendices

A Proofs of main theorems

A.1 Proof of Theorem 1

Denote the object function of (16) by $f(\Omega; \lambda_b, \Lambda_\Phi, \lambda_\Omega) = \ell(\Omega) + P_1(\Omega; \lambda_b, \Lambda_\Phi) + P_2(\Omega; \lambda_\Omega)$, where $P_1(\Omega; \lambda_b, \Lambda_\Phi) = \|\lambda_b \odot \mathbf{b}\|_1 + \|\Lambda_\Phi \odot \Phi\|_1$, and $P_2(\Omega; \lambda_\Omega) = \|(\lambda_\Omega \mathbf{1}^\top) \odot \Omega^\top\|_{2,1}$. Define $g(\Omega, \Omega'; \lambda_b, \Lambda_\Phi, \Lambda_\Omega) = \ell(\Omega') + \frac{\tau^2}{2} \|\Omega' - \Omega\|_F^2 + \langle \mathbf{X}^\top (\mathbf{y} - \mathbf{X} \mathbf{b}' - \bar{\mathbf{Z}} \text{vec}(\Phi')), \mathbf{b} - \mathbf{b}' \rangle + \langle \mathbf{Z}^\top \text{diag}\{\mathbf{y} - \mathbf{X} \mathbf{b}' - \bar{\mathbf{Z}} \text{vec}(\Phi')\} \mathbf{Z}, \Phi - \Phi' \rangle + P_1(\Omega; \lambda_b, \Lambda_\Phi) + P_2(\Omega; \lambda_\Omega)$. To simplify the notation, we sometimes write them as $f(\Omega)$ and $g(\Omega, \Omega')$ when there is no ambiguity.

Given Ω' , by simple algebra, the minimization of $g(\Omega, \Omega'; \lambda_b, \Lambda_\Phi, \lambda_\Omega)$ subject to $\Phi = \Phi^\top$ reduces to

$$\min_{\Omega \in \mathbb{R}^{(p+1) \times p}} \frac{1}{2} \|\Omega - \Xi\|_F^2 + P_1(\Omega; \lambda_b/\tau^2, \Lambda_\Phi/\tau^2) + P_2(\Omega; \lambda_\Omega/\tau^2) + \iota_{\{\Phi = \Phi^\top\}}, \quad (28)$$

where $\Xi = [\Xi_b, \Xi_\Phi]^\top$ with $\Xi_b = \mathbf{b}' + \mathbf{X}^\top(\mathbf{y} - \mathbf{X}\mathbf{b}' - \bar{\mathbf{Z}} \text{vec}(\Phi'))/\tau^2$ and $\Xi_\Phi = \Phi' + \mathbf{Z}^\top \text{diag}\{\mathbf{y} - \mathbf{X}\mathbf{b}' - \bar{\mathbf{Z}} \text{vec}(\Phi')\}\mathbf{Z}/\tau^2$, and $\iota_{\{\Phi = \Phi^\top\}} = 0$ if $\Phi = \Phi^\top$ and $+\infty$ otherwise, which is a lower semicontinuous convex function.

Lemma 1. *Let $\mathbf{P} \leftarrow \mathbf{0}$, $\mathbf{Q} \leftarrow \mathbf{0}$, $\Xi \leftarrow \Xi^0$. Consider the following iterative procedure:*

$$\begin{aligned} \Omega[:, k] &\leftarrow \vec{\Theta}_S(\Xi[:, k] + \mathbf{P}[:, k]; \lambda_\Omega[k]), \forall k : 1 \leq k \leq p, \\ \mathbf{P} &\leftarrow \mathbf{P} + \Xi - \Omega, \\ \Xi[1, :] &\leftarrow \Theta_S(\Omega[1, :] + \mathbf{Q}[1, :]; \lambda_b), \\ \Omega[2 : \text{end}, :] &\leftarrow (\Omega[2 : \text{end}, :] + \Omega[2 : \text{end}, :]^\top)/2, \\ \Xi[2 : \text{end}, :] &\leftarrow \Theta_S(\Omega[2 : \text{end}, :] + \mathbf{Q}[2 : \text{end}, :]; (\Lambda_\Phi + \Lambda_\Phi^\top)/2), \\ \mathbf{Q} &\leftarrow \mathbf{Q} + \Omega - \Xi, \end{aligned} \quad (29)$$

till convergence. Then both the sequence of Ξ and the sequence of Ω converge to the globally optimal solution to $\min_{\Omega \in \mathbb{R}^{(p+1) \times p}} \frac{1}{2} \|\Omega - \Xi^0\|_F^2 + \|\lambda_b \odot \Omega_b\|_1 + \|\Lambda_\Phi \odot \Omega_\Phi\|_1 + \|(\lambda_\Omega \mathbf{1}^\top) \odot \Omega^\top\|_{2,1}$ s.t. $\Omega_\Phi = \Omega_\Phi^\top$.

Now we prove the convergence of $\Omega^{(i)}$ in Algorithm 1. First, we notice the following fact.

Lemma 2. *Given any $\lambda \geq 0$, $\vec{\Theta}_S(\cdot; \lambda)$ is non-expansive, that is, $\|\vec{\Theta}(\mathbf{a}; \lambda) - \vec{\Theta}(\tilde{\mathbf{a}}; \lambda)\|_2 \leq \|\mathbf{a} - \tilde{\mathbf{a}}\|_2$ for all $\mathbf{a}, \tilde{\mathbf{a}} \in \mathbb{R}^p$.*

Therefore, the mapping of the iteration in Algorithm 1 is non-expansive. We use the tool provided by Opial (1967) for nonexpansive operators to prove the strict convergence of $\Omega^{(i)}$. First, the fix point set of the mapping is non-empty due to the convexity and the KKT conditions. The mapping is also asymptotically regular in the sense that for any starting point $\Omega^{(0)}$, $\|\Omega^{(i+1)} - \Omega^{(i)}\| \rightarrow 0$ as $i \rightarrow \infty$. This property is implied by the following lemma with $\tau > \|\check{\mathbf{Z}}\|_2$.

Lemma 3. For the sequence $\{\mathbf{\Omega}^{(i)}\}$ generated by Algorithm 1,

$$f(\mathbf{\Omega}^{(i)}; \boldsymbol{\lambda}_b, \mathbf{\Lambda}_\Phi, \boldsymbol{\lambda}_\Omega) - f(\mathbf{\Omega}^{(i+1)}; \boldsymbol{\lambda}_b, \mathbf{\Lambda}_\Phi, \boldsymbol{\lambda}_\Omega) \geq \frac{1}{2}(\tau^2 - \|\check{\mathbf{Z}}\|_2^2) \|\mathbf{\Omega}^{(i+1)} - \mathbf{\Omega}^{(i)}\|_F^2, \quad (30)$$

for any $i \geq 0$. Moreover, $\{\mathbf{\Omega}^{(i)}\}$ is uniformly bounded.

(30) actually does not require the optimality of $\mathbf{\Omega}^{(i+1)}$ in minimizing $g(\cdot, \mathbf{\Omega}^{(i)})$ (cf. (28)). From the proof of Lemma 3, any $\mathbf{\Omega}^{(i+1)}$ satisfying $g(\mathbf{\Omega}^{(i+1)}, \mathbf{\Omega}^{(i)}) \leq g(\mathbf{\Omega}^{(i)}, \mathbf{\Omega}^{(i)}) (= f(\mathbf{\Omega}^{(i)}))$ makes (30) hold. There is no need to run the inner loop till convergence. Moreover, to get the asymptotic regularity and the uniform boundedness of $\{\mathbf{\Omega}^{(i)}\}$, it suffices to terminate the inner loop when $g(\mathbf{\Omega}^{(i+1)}, \mathbf{\Omega}^{(i)}) \leq g(\mathbf{\Omega}^{(i)}, \mathbf{\Omega}^{(i)}) + \theta \|\mathbf{\Omega}^{(i+1)} - \mathbf{\Omega}^{(i)}\|_F^2/2$ or

$$\frac{g(\mathbf{\Omega}^{(i+1)}, \mathbf{\Omega}^{(i)}) - g(\mathbf{\Omega}^{(i)}, \mathbf{\Omega}^{(i)})}{\|\mathbf{\Omega}^{(i+1)} - \mathbf{\Omega}^{(i)}\|_F^2/2} \leq \theta,$$

where θ is a pre-specified constant satisfying $-\tau^2 \leq \theta < \tau^2 - \|\check{\mathbf{Z}}\|_2^2$. When $\mathbf{\Omega}^{(i+1)}$ exactly solves (28), $\theta = -\tau^2$ is satisfied.

With all Opial's conditions satisfied, the sequence $\{\mathbf{\Omega}^{(i)}\}$ has a unique limit point $\mathbf{\Omega}^* = [\mathbf{b}^*, \mathbf{\Phi}^{*\top}]^\top$, and $\mathbf{\Omega}^*$ is a fixed point of Algorithm 1.

Next we prove the optimality of any fixed point $\mathbf{\Omega}^*$ regarding problem (16). By Lemma 1, $\mathbf{\Omega}^*$ satisfies the Karush-Kuhn-Tucker (KKT) conditions of (28):

$$\begin{cases} \mathbf{0} \in \mathbf{\Omega}^* - \Xi + \partial P_1(\mathbf{\Omega}; \boldsymbol{\lambda}_b/\tau^2, \mathbf{\Lambda}_\Phi/\tau^2) + \partial P_2(\mathbf{\Omega}; \mathbf{\Lambda}_\Omega/\tau^2) + \partial h(\mathbf{\Phi}^*), \\ \mathbf{\Phi}^* - \mathbf{\Phi}^{*\top} = \mathbf{0}, \end{cases}$$

where $h(\mathbf{\Phi}; \mathbf{L}) = \langle \mathbf{L}, \mathbf{\Phi} - \mathbf{\Phi}^\top \rangle$. With the fixed point property established above, substituting Ξ with $\Xi_b = \mathbf{b}^* + \mathbf{X}^\top(\mathbf{y} - \mathbf{X}\mathbf{b}^* - \bar{\mathbf{Z}} \text{vec}(\mathbf{\Phi}^*))/\tau^2$ and $\Xi_\Phi = \mathbf{\Phi}^* + \mathbf{Z}^\top \text{diag}\{\mathbf{y} - \mathbf{X}\mathbf{b}^* - \bar{\mathbf{Z}} \text{vec}(\mathbf{\Phi}^*)\}\mathbf{Z}/\tau^2$, we have

$$\begin{cases} \mathbf{0} \in \mathbf{X}^\top(\mathbf{y} - \mathbf{X}\mathbf{b}^* - \bar{\mathbf{Z}} \text{vec}(\mathbf{\Phi}^*)) + \partial_b P_1(\mathbf{\Omega}; \boldsymbol{\lambda}_b, \mathbf{\Lambda}_\Phi) + \partial_b P_2(\mathbf{\Omega}; \mathbf{\Lambda}_\Omega), \\ \mathbf{0} \in \mathbf{Z}^\top \text{diag}\{\mathbf{y} - \mathbf{X}\mathbf{b}^* - \bar{\mathbf{Z}} \text{vec}(\mathbf{\Phi}^*)\}\mathbf{Z} + \partial_\Phi(P_1(\mathbf{\Omega}; \boldsymbol{\lambda}_b, \mathbf{\Lambda}_\Phi) + P_2(\mathbf{\Omega}; \mathbf{\Lambda}_\Omega)) \\ \quad + \partial h(\mathbf{\Phi}^*; \tau^2 \mathbf{L}), \\ \mathbf{0} = \mathbf{\Phi}^* - \mathbf{\Phi}^{*\top}, \end{cases}$$

which are exactly the KKT conditions for the convex problem (16). Hence $\mathbf{\Omega}^*$ is a global minimizer of problem (16).

Similarly, the global optimality of $\lim_{i \rightarrow \infty} \mathbf{\Omega}^{(i)} = \mathbf{\Omega}^*$ does not require solving (28) exactly. In the following, we use $\|\mathbf{A}\|_{\Sigma}$ to denote $\{\text{tr}(\mathbf{A}^\top \Sigma \mathbf{A})\}^{1/2}$. Let $\mathbf{\Omega}^o$ be an arbitrary optimal solution to (16). From the proof of Lemma 3, the optimal $\mathbf{\Omega}^{(i+1)}$ satisfies $g(\mathbf{\Omega}^{(i+1)}, \mathbf{\Omega}^{(i)}) + \frac{\tau^2}{2} \|\mathbf{\Omega}^{(i+1)} - \mathbf{\Omega}^o\|_F^2 \leq g(\mathbf{\Omega}^o, \mathbf{\Omega}^{(i)})$. We can relax it to

$$g(\mathbf{\Omega}^{(i+1)}, \mathbf{\Omega}^{(i)}) + \frac{1}{2} \|\mathbf{\Omega}^{(i+1)} - \mathbf{\Omega}^o\|_{\tau^2 \mathbf{I} - \theta \mathbf{X}^\top \mathbf{X}}^2 \leq g(\mathbf{\Omega}^o, \mathbf{\Omega}^{(i)})$$

for some $\theta : 0 \leq \theta \leq 1$. Then by the construction of g , we obtain

$$\begin{aligned} f(\mathbf{\Omega}^{(i+1)}) - f(\mathbf{\Omega}^o) &\leq \frac{1}{2} \|\mathbf{\Omega}^{(i)} - \mathbf{\Omega}^o\|_{\tau^2 \mathbf{I} - \mathbf{X}^\top \mathbf{X}}^2 - \frac{1}{2} \|\mathbf{\Omega}^{(i+1)} - \mathbf{\Omega}^o\|_{\tau^2 \mathbf{I} - \theta \mathbf{X}^\top \mathbf{X}}^2 \\ &\quad - \frac{1}{2} \|\mathbf{\Omega}^{(i+1)} - \mathbf{\Omega}^{(i)}\|_{\tau^2 \mathbf{I} - \mathbf{X}^\top \mathbf{X}}^2 \\ &\leq \|\mathbf{\Omega}^{(i)} - \mathbf{\Omega}^o\|_{\tau^2 \mathbf{I} - \mathbf{X}^\top \mathbf{X}}^2 - \frac{1}{2} \|\mathbf{\Omega}^{(i+1)} - \mathbf{\Omega}^o\|_{\tau^2 \mathbf{I} - \theta \mathbf{X}^\top \mathbf{X}}^2 \\ &\leq \frac{1}{2} \|\mathbf{\Omega}^{(i)} - \mathbf{\Omega}^o\|_{\tau^2 \mathbf{I} - \theta \mathbf{X}^\top \mathbf{X}}^2 - \frac{1}{2} \|\mathbf{\Omega}^{(i+1)} - \mathbf{\Omega}^o\|_{\tau^2 \mathbf{I} - \theta \mathbf{X}^\top \mathbf{X}}^2. \end{aligned}$$

It follows that

$$\bar{f}_{i+1} - f(\mathbf{\Omega}^o) \leq \frac{1}{2(i+1)} \|\mathbf{\Omega}^{(0)} - \mathbf{\Omega}^o\|_{\tau^2 \mathbf{I} - \theta \mathbf{X}^\top \mathbf{X}}^2, \quad (31)$$

where $\bar{f}_{i+1} = \frac{1}{i+1} \sum_{k=1}^{i+1} f(\mathbf{\Omega}^{(k)})$. Letting $i \rightarrow \infty$, we get $f(\mathbf{\Omega}^*) = f(\lim_{i \rightarrow \infty} \mathbf{\Omega}^{(i)}) = \lim_{i \rightarrow \infty} \bar{f}_i = f(\mathbf{\Omega}^o)$, and so $\mathbf{\Omega}^*$ must also be a global minimizer of (16).

Proof of Lemma 1 First, we consider

$$\min_{\mathbf{\Omega}=[\mathbf{b}, \mathbf{\Phi}]^\top} \frac{1}{2} \|\mathbf{\Omega} - \mathbf{\Xi}^0\|_F^2 + \|\mathbf{\lambda}_b \odot \mathbf{b}\|_1 + \|\mathbf{\Lambda}_\Phi \odot \mathbf{\Phi}\|_1 + \iota_{\{\mathbf{\Phi}=\mathbf{\Phi}^\top\}}.$$

Let $\mathbf{\Phi}^0 = \mathbf{\Xi}_\Phi^0$, $\mathbf{b}^0 = \mathbf{\Xi}_b^0$. By some simple algebra, $\|\mathbf{\Phi} - \mathbf{\Phi}^0\|_F^2 = \|\mathbf{\Phi} - (\mathbf{\Phi}^0 + \mathbf{\Phi}^{0\top})/2\|_F^2 + C(\mathbf{\Phi}^0)$, where the last term does not depend on $\mathbf{\Phi}$. Based on Lemma 1 of She (2012), it is not difficult to show that the global minimizer is given by $\mathbf{b}_o = \Theta_S(\mathbf{b}^0; \mathbf{\lambda}_b)$, $\mathbf{\Phi}_o = \Theta_S((\mathbf{\Phi}^0 + \mathbf{\Phi}^{0\top})/2; (\mathbf{\Lambda}_\Phi + \mathbf{\Lambda}_\Phi^\top)/2)$.

Similarly, the globally optimal solution to $\min_{\mathbf{\Omega}=[\mathbf{b}, \mathbf{\Phi}]^\top} \frac{1}{2} \|\mathbf{\Omega} - \mathbf{\Xi}^0\|_F^2 + \|(\mathbf{\lambda}_\Omega \mathbf{1}^\top) \odot \mathbf{\Omega}^\top\|_{2,1}$ is given by $\mathbf{\Omega}_o$ with $\mathbf{\Omega}_o[:, k] = \bar{\Theta}_S(\mathbf{\Xi}^0[:, k]; \mathbf{\lambda}_\Omega[k])$, $1 \leq k \leq p$. Applying Theorem 3.2 and 3.3 in Bauschke and Combettes (2008) guarantees the strict convergence of the iterates and the global optimality of the limit point. \square

Proof of Lemma 2 The conclusion follows from the fact that $\bar{\Theta}$ is a proximity operator associated with a convex function. \square

Proof of Lemma 3 From the optimality of $\Omega^{(i+1)}$ and the convexity of the penalties in (28), it is easy to see that for any Ω ,

$$g(\Omega^{(i+1)}, \Omega^{(i)}; \lambda_b, \Lambda_\Phi, \Lambda_\Omega) + \frac{\tau^2}{2} \|\Omega^{(i+1)} - \Omega\|_F^2 \leq g(\Omega, \Omega^{(i)}; \lambda_b, \Lambda_\Phi, \lambda_\Omega). \quad (32)$$

It follows from the construction of g that

$$g(\Omega^{(i+1)}, \Omega^{(i)}; \lambda_b, \Lambda_\Phi, \Lambda_\Omega) \leq g(\Omega^{(i)}, \Omega^{(i)}; \lambda_b, \Lambda_\Phi, \lambda_\Omega) = f(\Omega^{(i)}; \lambda_b, \Lambda_\Phi, \lambda_\Omega).$$

On the other hand, noticing that the gradient of l with respect to Φ is $\mathbf{Z}^\top \text{diag}\{\mathbf{y} - \mathbf{X}\mathbf{b}' - \bar{\mathbf{Z}} \text{vec}(\Phi')\}\mathbf{Z}$, Taylor expansion yields

$$g(\Omega, \Omega'; \lambda_b, \Lambda_\Phi, \Lambda_\Omega) \geq f(\Omega; \lambda_b, \Lambda_\Phi, \lambda_\Omega) + \frac{(\tau^2 - \|\check{\mathbf{Z}}\|_2^2)}{2} \|\Omega' - \Omega\|_F^2$$

for any Ω, Ω' .

Furthermore, $\|(\lambda_\Omega \mathbf{1}^\top) \odot \Omega^{(i)\top}\|_{2,1} \leq f(\Omega^{(i)}; \lambda_b, \Lambda_\Phi) \leq f(\Omega^{(0)}; \lambda_b, \Lambda_\Phi)$. Because $\lambda_\Omega > 0$, the conclusion follows. \square

A.2 Proof of Theorem 2

In this proof, we use C, c to denote universal constants. They are not necessarily the same at each occurrence. Throughout the proof, $P(\mathbf{b}, \Phi; \lambda_1, \lambda_2)$ is short for $\lambda_1 \|\Phi\|_1 + \lambda_2 \|\Omega^\top\|_{2,1}$, λ'_1 is short for $\lambda_1 \|\check{\mathbf{X}}\|_2$, and λ'_2 is short for $\lambda_2 \|\check{\mathbf{X}}\|_2$. We prove the theorem under a less restrictive condition.

First, because $\hat{\Omega}$ is a global minimizer of (8) and P is convex, for any $(\mathbf{b}, \Phi) \in \mathbb{R}^p \times \mathbb{R}^{p \times p}$ with Φ symmetric (not necessarily satisfying SH), we have

$$\ell(\hat{\mathbf{b}}, \hat{\Phi}) + P(\hat{\mathbf{b}}, \hat{\Phi}; \lambda'_1, \lambda'_2) + \frac{1}{2} M(\mathbf{b} - \hat{\mathbf{b}}, \Phi - \hat{\Phi}) \leq \ell(\mathbf{b}, \Phi) + P(\mathbf{b}, \Phi; \lambda'_1, \lambda'_2).$$

Based on the model assumption, it is easy to see that

$$\begin{aligned} & \frac{1}{2} M(\mathbf{b}^* - \hat{\mathbf{b}}, \Phi^* - \hat{\Phi}) + P(\hat{\mathbf{b}}, \hat{\Phi}; \lambda'_1, \lambda'_2) + \frac{1}{2} M(\Delta_b, \Delta_\Phi) \\ & \leq \frac{1}{2} M(\mathbf{b}^* - \mathbf{b}, \Phi^* - \Phi) + P(\mathbf{b}, \Phi; \lambda'_1, \lambda'_2) + \langle \epsilon, \mathbf{X} \Delta_b + \text{diag}(\mathbf{X} \Delta_\Phi \mathbf{X}^\top) \rangle, \end{aligned}$$

where $\Delta_b = \hat{\mathbf{b}} - \mathbf{b}$, $\Delta_\Phi = \hat{\Phi} - \Phi$.

Introduce $\check{\Delta} = [\Delta_b^\top, (\text{vec}(\Delta_\Phi))^\top]^\top$ and $\tilde{\mathcal{J}} = \mathcal{J} \cup \hat{\mathcal{J}}$, where $\mathcal{J} := \{j : \text{vec}(\Omega)_j \neq 0\}$. and $\hat{\mathcal{J}} := \{j : \text{vec}(\hat{\Omega})_j \neq 0\}$. Denote by $\mathcal{P}_{\tilde{\mathcal{J}}}$ the orthogonal projection matrix onto the column space of $\check{\mathbf{X}}_{\tilde{\mathcal{J}}}$. Then

$$\begin{aligned} & \langle \epsilon, \mathbf{X}\Delta_b + \text{diag}(\mathbf{X}\Delta_\Phi\mathbf{X}^\top) \rangle \\ &= \langle \epsilon, \check{\mathbf{X}}\check{\Delta} \rangle \\ &= \langle \epsilon, \mathcal{P}_{\tilde{\mathcal{J}}} \check{\mathbf{X}}_{\tilde{\mathcal{J}}} \check{\Delta}_{\tilde{\mathcal{J}}} \rangle \\ &= \langle \mathcal{P}_{\tilde{\mathcal{J}}} \epsilon, \mathbf{X}\Delta_b + \text{diag}(\mathbf{X}\Delta_\Phi\mathbf{X}^\top) \rangle \\ &\leq \frac{1}{2a_1} M(\Delta_b, \Delta_\Phi) + \frac{a_1}{2} \|\mathcal{P}_{\tilde{\mathcal{J}}} \epsilon\|_2^2, \end{aligned}$$

for any $a_1 > 0$. Representing $\|\hat{\Omega}\|_{2,1}$ as $\|(\hat{\Omega}_{\mathcal{J}_G})^\top\|_{2,1} + \|(\hat{\Omega}_{\mathcal{J}_G^c})^\top\|_{2,1}$ and $\|\hat{\Phi}\|_1$ as $\|\hat{\Phi}_{\mathcal{J}_e}\|_1 + \|\hat{\Phi}_{\mathcal{J}_e^c}\|_1$, respectively, and applying the triangle inequalities of $\|\cdot\|_1$ and $\|\cdot\|_{2,1}$, we have

$$\begin{aligned} & \frac{1}{2} M(\mathbf{b}^* - \hat{\mathbf{b}}, \Phi^* - \hat{\Phi}) + \lambda'_1 \|(\Delta_\Phi)_{\mathcal{J}_e}\|_1 + \lambda'_2 \|(\Delta_{\mathcal{J}_G^c})^\top\|_{2,1} + \frac{1}{2} M(\Delta_b, \Delta_\Phi) \\ &\leq \frac{1}{2} M(\mathbf{b}^* - \mathbf{b}, \Phi^* - \Phi) + \frac{1}{2a_1} M(\Delta_b, \Delta_\Phi) + \lambda'_1 \|(\Delta_\Phi)_{\mathcal{J}_e}\|_1 + \lambda'_2 \|(\Delta_{\mathcal{J}_G})^\top\|_{2,1} \\ &\quad + \frac{a_1}{2} \|\mathcal{P}_{\tilde{\mathcal{J}}} \epsilon\|_2^2. \end{aligned} \tag{33}$$

How to treat the key term $\|\mathcal{P}_{\tilde{\mathcal{J}}} \epsilon\|_2^2$ is nontrivial. We derive two inequalities from statistical analysis and computational analysis.

First, we bound $\|\mathcal{P}_{\tilde{\mathcal{J}}} \epsilon\|_2^2$ with some cardinality measures. The goal of this step is to show for a large constant L , $\sup_{\mathcal{J} \subset [p(p+1)]} \{\|\mathcal{P}_{\tilde{\mathcal{J}}} \epsilon\|_2^2 - LP_o(J_G(\mathcal{J}), J_e(\mathcal{J}))\}$ is negative with high probability, where $P_o(J_G, J_e) = \sigma^2 \{J_G \log(ep/J_G) + J_e \log(epJ_G/J_e)\}$. (If the reference signal (\mathbf{b}, Φ) satisfies SH, so does $\tilde{\mathcal{J}}$, for which P_o can be strengthened to $\sigma^2 \{J_G \log(ep/J_G) + J_e \log(eJ_G^2/J_e)\}$ following the lines of the proof of Lemma 4.)

For notational convenience, we extend the J -measures defined in (2) to an index set: Given any $\mathcal{J} \subset [p(p+1)]$, with $\Omega = [\mathbf{b}, \Phi^\top]^\top$ a binary matrix satisfying $(\text{vec}(\Omega))_{\mathcal{J}} = 1$ and $(\text{vec}(\Omega))_{\mathcal{J}^c} = 0$, define $\mathcal{J}_G(\mathcal{J}) = \mathcal{J}_G(\Omega)$, $\mathcal{J}_e(\mathcal{J}) = \mathcal{J}_e(\Phi)$. Given $\mathcal{J} \subset [p(p+1)]$, $\mathcal{J}_G(\mathcal{J}) \subset [p]$ and $\mathcal{J}_e(\mathcal{J}) \subset [p^2]$.

Lemma 4. *Given $1 \leq J_G \leq p$ and $0 \leq J_e \leq p^2$, for any $t \geq 0$,*

$$\mathbb{P} \left(\sup_{|\mathcal{J}_G|=J_G, |\mathcal{J}_e|=J_e} \|\mathcal{P}_{\mathcal{J}} \boldsymbol{\varepsilon}\|_2^2 \geq L_0 P_o(J_G, J_e) + t\sigma^2 \right) \leq C \exp(-ct), \quad (34)$$

where $L_0, C, c > 0$ are universal constants. Here, we omit the dependence of \mathcal{J}_G and \mathcal{J}_e on \mathcal{J} for brevity.

Define

$$R_{J_G, J_e} = \left(\sup_{|\mathcal{J}_G|=J_G, |\mathcal{J}_e|=J_e} \|\mathcal{P}_{\mathcal{J}} \boldsymbol{\varepsilon}\|_2^2 - L\sigma^2 \{J_e \log(ep) + J_G \log(ep)\} \right)_+,$$

and $R = \sup_{1 \leq J_G \leq p, 0 \leq J_e \leq p^2} R_{J_G, J_e}$ (when $J_G(\mathcal{J}) = 0$, $\mathcal{P}_{\mathcal{J}} \boldsymbol{\varepsilon} = \mathbf{0}$). Then

$$\begin{aligned} \|\mathcal{P}_{\tilde{\mathcal{J}}} \boldsymbol{\varepsilon}\|_2^2 &\leq L\sigma^2 \{\tilde{J}_e \log(ep) + \tilde{J}_G \log(ep)\} + R \\ &\leq L\sigma^2 (J_e + J_G) \log(ep) + L\sigma^2 \{\hat{J}_e \log(ep) + \hat{J}_G \log(ep)\} + R. \end{aligned} \quad (35)$$

Set $L > L_0$. Then, by Lemma 4 and $J_e \geq J_G$,

$$\begin{aligned} &\mathbb{P}(R \geq t\sigma) \\ &\leq \sum_{J_G=1}^p \sum_{J_e=0}^{p^2} \mathbb{P}(R_{J_G, J_e} \geq t\sigma) \\ &\leq \sum_{J_G=1}^p \sum_{J_e=0}^{p^2} C \exp(-ct) \exp \{-c(L - L_0)(J_e \log(ep) + J_G \log(ep))\} \\ &\leq C' \exp(-ct^2), \end{aligned}$$

from which it follows that $\mathbb{E}R \leq C\sigma^2$. It is also easy to see that for sufficiently large L , $R \leq 0$ occurs with probability at least $1 - Cp^{-cL}$.

Next, we derive an inequality based on the computational optimality of $(\hat{\mathbf{b}}, \hat{\Phi})$. Due to the convexity of the problem, $(\hat{\mathbf{b}}, \hat{\Phi})$ is a stationary point of

$$\begin{aligned} &\frac{1}{2} \|\mathbf{y} - \mathbf{X}\mathbf{b} - \text{diag}(\mathbf{X}\Phi\mathbf{X}^\top)\|_2^2 + \lambda'_1 \|\Phi\|_1 \\ &\quad + \lambda'_2 \|\Omega^\top\|_{2,1} + \sum_{j < k} l_{j,k}(\phi_{j,k} - \phi_{k,j}), \end{aligned}$$

where $l_{j,k}$ are Lagrangian multipliers. By the KKT conditions, for any $j \in \hat{\mathcal{J}}_G$, \hat{b}_j is nonzero (with probability 1) and so satisfies

$$\lambda'_2 \frac{\hat{b}_j}{\|\hat{\mathbf{\Omega}}_j\|_2} = -\mathbf{x}_j^\top (\check{\mathbf{X}} \text{vec}(\hat{\mathbf{\Omega}}) - \mathbf{y}). \quad (36)$$

Similarly, for any $j' \in \hat{\mathcal{J}}_e$, letting $j_c = \lceil \frac{j'}{p} \rceil$, $j_r = j' - (j_c - 1)p$, $\hat{\phi}_{j_r, j_c}$ satisfies $\lambda'_1 \text{sgn}(\hat{\phi}_{j_r, j_c}) + \lambda'_2 \frac{\hat{\phi}_{j_r, j_c}}{\|\hat{\mathbf{\Omega}}_{j_c}\|_2} + l_{j_r, j_c} \text{sgn}(j_r - j_c) = -\bar{\mathbf{X}}_{j'}^\top (\check{\mathbf{X}} \text{vec}(\hat{\mathbf{\Omega}}) - \mathbf{y})$, where $\text{sgn}(t) = 1, 0, -1$ for $t > 0, = 0, < 0$, respectively. Adding the equations for $\hat{\phi}_{j_r, j_c}$ and $\hat{\phi}_{j_c, j_r}$ can cancel l_{j_r, j_c} . So when $j_r \neq j_c$, we have

$$2\lambda'_1 \text{sgn}(\hat{\phi}_{j_r, j_c}) + \lambda'_2 \frac{\hat{\phi}_{j_r, j_c}}{\|\hat{\mathbf{\Omega}}_{j_c}\|_2} + \lambda'_2 \frac{\hat{\phi}_{j_c, j_r}}{\|\hat{\mathbf{\Omega}}_{j_r}\|_2} = -2\bar{\mathbf{X}}_{j'}^\top (\check{\mathbf{X}} \text{vec}(\hat{\mathbf{\Omega}}) - \mathbf{y}), \quad (37)$$

and when $j_c = j_r$, we have

$$\lambda'_1 \text{sgn}(\hat{\phi}_{j_r, j_c}) + \lambda'_2 \frac{\hat{\phi}_{j_r, j_c}}{\|\hat{\mathbf{\Omega}}_{j_c}\|_2} = -\bar{\mathbf{X}}_{j'}^\top (\check{\mathbf{X}} \text{vec}(\hat{\mathbf{\Omega}}) - \mathbf{y}). \quad (38)$$

Notice that $\text{sgn}(\hat{\phi}_{j_r, j_c})\hat{\phi}_{j_r, j_c}$ is always non-negative. Squaring both sides of (36), (37) and (38) and summing over all $j \in \hat{\mathcal{J}}_G$ and $j' \in \hat{\mathcal{J}}_e$, we obtain

$$\lambda_1^2 \hat{J}_e + \lambda_2^2 \sum_{j \in \hat{\mathcal{J}}_G} \frac{\hat{b}_j^2 + \sum_{k: \hat{\phi}_{k,j} \neq 0} \hat{\phi}_{k,j}^2}{\|\hat{\mathbf{\Omega}}_j\|_2^2} \leq 2\|\check{\mathbf{X}}_{\hat{\mathcal{J}}}^\top (\check{\mathbf{X}} \text{vec}(\hat{\mathbf{\Omega}}) - \mathbf{y})\|_2^2,$$

and so $\lambda_1^2 \hat{J}_e + \lambda_2^2 \hat{J}_G \leq 2\|\check{\mathbf{X}}_{\hat{\mathcal{J}}}^\top (\check{\mathbf{X}} \text{vec}(\hat{\mathbf{\Omega}}) - \mathbf{y})\|_2^2$. It follows from $\hat{\mathcal{J}} \subset \tilde{\mathcal{J}}$ and the Cauchy-Schwarz inequality that

$$\lambda_1^2 \hat{J}_e + \lambda_2^2 \hat{J}_G \leq 3M(\mathbf{b}^* - \hat{\mathbf{b}}, \mathbf{\Phi}^* - \hat{\mathbf{\Phi}}) + 6\|\mathcal{P}_{\tilde{\mathcal{J}}} \boldsymbol{\varepsilon}\|_2^2. \quad (39)$$

Now, combining the optimization inequality (39) and the statistical inequality (35) can yield a bound of $\|\mathcal{P}_{\tilde{\mathcal{J}}} \boldsymbol{\varepsilon}\|_2^2$. In fact, with $\lambda^\circ = \sigma \sqrt{\log(ep)}$, $\lambda_1 = A_1 \lambda^\circ$, $\lambda_2 = A_2 \lambda^\circ$, $A_i \geq A$, $A^2 = A_0 L$ for some $A_0 > 0$, we get

$$(1 - \frac{6}{A_0})\|\mathcal{P}_{\tilde{\mathcal{J}}} \boldsymbol{\varepsilon}\|_2^2 \leq L\sigma^2(J_e + J_G) \log(ep) + \frac{3}{A_0}M(\mathbf{b}^* - \hat{\mathbf{b}}, \mathbf{\Phi}^* - \hat{\mathbf{\Phi}}) + R.$$

Plugging this into (33) gives

$$\begin{aligned}
& \frac{1}{2}M(\mathbf{b}^* - \hat{\mathbf{b}}, \Phi^* - \hat{\Phi}) + A\lambda^o \|\check{\mathbf{X}}\|_2 \|(\Delta_\Phi)_{\mathcal{J}_e}\|_1 + A\lambda^o \|\check{\mathbf{X}}\|_2 \|(\Delta_{\mathcal{J}_G^c})^\top\|_{2,1} \\
& \leq \frac{1}{2} \left(1 + \frac{3a_1}{A_0 - 6}\right) M(\mathbf{b}^* - \mathbf{b}, \Phi^* - \Phi) + \frac{1}{2} \left(\frac{1}{a_1} - 1\right) M(\Delta_b, \Delta_\Phi) \\
& \quad + A\lambda^o \|\check{\mathbf{X}}\|_2 \|(\Delta_\Phi)_{\mathcal{J}_e}\|_1 + A\lambda^o \|\check{\mathbf{X}}\|_2 \|(\Delta_{\mathcal{J}_G})^\top\|_{2,1} + \frac{a_1 L}{2(1 - 6/A_0)} \sigma^2(J_e + J_G) \log(ep) \\
& \quad + \frac{a_1}{2(1 - 6/A_0)} R.
\end{aligned}$$

Assume the following condition holds

$$\begin{aligned}
& \|\check{\mathbf{X}}\|_2 (\|(\Delta_\Phi)_{\mathcal{J}_e}\|_1 + \|(\Delta_{\mathcal{J}_G})^\top\|_{2,1}) \\
& \leq \|\check{\mathbf{X}}\|_2 (\|(\Delta_\Phi)_{\mathcal{J}_e}\|_1 + \|(\Delta_{\mathcal{J}_G^c})^\top\|_{2,1}) + K \sqrt{J_e + J_G} \|\mathbf{X} \Delta_b + \bar{\mathbf{X}} \text{vec}(\Delta_\Phi)\|_2,
\end{aligned} \tag{40}$$

for some K large enough. Then we have

$$\begin{aligned}
& \frac{1}{2}M(\mathbf{b}^* - \hat{\mathbf{b}}, \Phi^* - \hat{\Phi}) \\
& \leq \frac{1}{2} \left(1 + \frac{3a_1}{A_0 - 6}\right) M(\mathbf{b}^* - \mathbf{b}, \Phi^* - \Phi) + \frac{1}{2} \left(\frac{1}{a_1} - 1\right) M(\Delta_b, \Delta_\Phi) + \frac{a_1 R}{2(1 - 6/A_0)} \\
& \quad + A\lambda^o K \sqrt{J_e + J_G} \|\mathbf{X} \Delta_b + \bar{\mathbf{X}} \text{vec}(\Delta_\Phi)\|_2 + \frac{a_1 L}{2(1 - 6/A_0)} \sigma^2(J_e + J_G) \log(ep) \\
& \leq \frac{1}{2} \left(1 + \frac{3a_1}{A_0 - 6}\right) M(\mathbf{b}^* - \mathbf{b}, \Phi^* - \Phi) + \frac{1}{2} \left(\frac{1}{a_1} + \frac{1}{a_2} - 1\right) M(\Delta_b, \Delta_\Phi) \\
& \quad + \left\{ \frac{a_1}{2(1 - 6/A_0)} + \frac{a_2 A_0}{2} K^2 \right\} L \sigma^2(J_e + J_G) \log(ep) + \frac{a_1}{2(1 - 6/A_0)} R.
\end{aligned}$$

Hence we get $\mathbb{E}M(\mathbf{b}^* - \hat{\mathbf{b}}, \Phi^* - \hat{\Phi}) \lesssim M(\mathbf{b}^* - \mathbf{b}, \Phi^* - \Phi) + (1 \vee K^2) \sigma^2(J_e + J_G) \sigma^2 \log(ep) + \sigma^2$ if we choose the constants a_1, a_2, A_0 satisfying, say, $a_1 = 2, a_2 = 2, A_0 \geq 7$.

To complete the proof, we show that $\mathcal{A}(\mathcal{J}_e, \mathcal{J}_G, \vartheta, \kappa)$ with $\kappa > 0$ and $\vartheta \geq 0$ implies (40). Consider two cases. The case $(1 + \vartheta)(\|(\Delta_\Phi)_{\mathcal{J}_e}\|_1 + \|(\Delta_{\mathcal{J}_G})^\top\|_{2,1}) \leq \|(\Delta_\Phi)_{\mathcal{J}_e}\|_1 + \|(\Delta_{\mathcal{J}_G^c})^\top\|_{2,1}$ is trivial. Suppose the reverse

inequality holds. Then

$$\begin{aligned}
& \|\check{\mathbf{X}}\|_2 (\|(\Delta_\Phi)_{\mathcal{J}_e}\|_1 + \|(\Delta_{\mathcal{J}_G})^\top\|_{2,1}) \\
& \leq \|\check{\mathbf{X}}\|_2 (J_e + J_G)^{1/2} (\|(\Delta_\Phi)_{\mathcal{J}_e}\|_2^2 + \|(\Delta_{\mathcal{J}_G})^\top\|_F^2)^{1/2} \\
& \leq (J_e + J_G)^{1/2} \|\mathbf{X}\Delta\|_F / \kappa^{1/2},
\end{aligned}$$

and so (40) holds with $K = 1/\kappa^{1/2}$.

Finally, we can also plug (35) into (39), resulting in

$$\lambda_1^2 \mathbb{E}[\hat{J}_e] + \lambda_2^2 \mathbb{E}[\hat{J}_G] \lesssim M(\mathbf{b}^* - \mathbf{b}, \Phi^* - \Phi) + \lambda_1^2 J_e + \lambda_2^2 J_G + \sigma^2, \quad (41)$$

and so (22) follows.

The proof for Theorem 2' follows the same lines; the details are omitted.

Proof of Lemma 4 First, notice that for fixed \mathcal{J} , $\|\mathcal{P}_{\mathcal{J}}\varepsilon\|_2^2/\sigma^2 \sim \chi^2(D)$ with $D \leq J$. The standard tail bound for $X \sim \chi^2(D)$ gives $\mathbb{P}(X - D \geq t) \leq \exp(-t^2/(4(D+t)))$, $\forall t \geq 0$ (Laurent and Massart, 2000, Lemma 1).

It is easy to see that

$$\sup_{|\mathcal{J}_G|=J_G, |\mathcal{J}_e|=J_e} \|\mathcal{P}_{\mathcal{J}}\varepsilon\|_2^2 \leq \sup_{J^{11}+J^{10}+J^{01}=J_G} \sup_{\mathcal{Q}(J^{11}, J^{10}, J^{01})} \sup_{|\mathcal{J}_e|=J_e, \mathcal{J}_e \subset \mathcal{J}_G^{(e)}} \|\mathcal{P}_{\mathcal{J}}\varepsilon\|_2^2,$$

where $\mathcal{Q}(J^{11}, J^{10}, J^{01}) = \{(\mathcal{J}^{11}, \mathcal{J}^{10}, \mathcal{J}^{01}) \subset [p]^3 : |\mathcal{J}^{11}| = J^{11}, |\mathcal{J}^{10}| = J^{10}, |\mathcal{J}^{01}| = J^{01}, \mathcal{J}^{11} \cap \mathcal{J}^{10} = \mathcal{J}^{10} \cap \mathcal{J}^{01} = \mathcal{J}^{11} \cap \mathcal{J}^{01} = \emptyset\}$ and $\mathcal{J}_G^{(e)}$ denotes the index set $\{1 \leq j \leq p^2 : \lceil \frac{j}{p} \rceil \in \mathcal{J}_G\}$. $\mathcal{Q}(J^{11}, J^{10}, J^{01})$ contains

$$\binom{p}{J^{11}, J^{10}, J^{01}, p - J^{11} - J^{10} - J^{01}}$$

many elements. This multinomial coefficient is bounded by $\binom{p}{J_G}$, and by Stirling's approximation, $\log \binom{p}{J_G} \leq J_G \log(ep/J_G)$. Similarly, $\sup_{|\mathcal{J}_e|=J_e, \mathcal{J}_e \subset \mathcal{J}_G^{(e)}}$ and $\sup_{J^{11}+J^{10}+J^{01}=J_G}$ involve $\binom{(J^{11}+J^{01})p}{J_e}$ and $\binom{J_G+2}{2}$ terms, respectively, and we have

$$\log \binom{(J^{11}+J^{01})p}{J_e} \leq J_e \log(ep/J_e), \quad \log \binom{J_G+2}{2} \leq C \log(eJ_G).$$

Applying the union bound gives the desired result. \square

A.3 Proof of Theorem 3

The proof is based on the general reduction scheme in Chapter 2 of Tsybakov (2009). The key is to design proper least favorable signals in different situations. We consider two cases.

(i) $J_G \log(ep/J_G) \leq J_e \log(eJ_G^2/J_e)$. Define a signal subclass

$$\mathcal{B}^1(J_G, J_e) = \left\{ \mathbf{\Omega} = [\mathbf{b}, \mathbf{\Phi}]^\top : b_j = 1 \text{ if } j \in [J_G] \text{ and } 0 \text{ if } J_G < j \leq p, \right. \\ \left. \phi_{k,j} = 0 \text{ or } \gamma R \text{ if } (k, j) \in [J_G] \times [J_G] \text{ and } 0 \text{ otherwise, } \mathbf{\Phi} = \mathbf{\Phi}^\top, J_e(\mathbf{\Phi}) \leq J_e \right\},$$

where $R = \frac{\sigma}{\kappa^{1/2} \|\check{\mathbf{X}}\|_2} (\log(eJ_G^2/J_e))^{1/2}$ and $\gamma > 0$ is a small constant to be chosen later. Because $b_j = 1$ ($1 \leq j \leq J_G$), any $\mathbf{\Omega} \in \mathcal{B}^1(J_G, J_e)$ satisfies SH, and thus $\mathcal{B}^1(J_G, J_e) \subset \text{SH}(J_G, J_e)$. By Stirling's approximation,

$$\log |\mathcal{B}^1(J_G, J_e)| = \log \binom{J_G^2}{J_e} \geq J_e \log(J_G^2/J_e) \geq c J_e \log(eJ_G^2/J_e)$$

for some universal constant c .

Let $\rho(\mathbf{\Omega}_1, \mathbf{\Omega}_2) = \|\mathbf{\Omega}_1 - \mathbf{\Omega}_2\|_0$ be the Hamming distance. By Lemma 8.3 in Rigollet and Tsybakov (2011), there exists a subset $\mathcal{B}^{10}(J_G, J_e) \subset \mathcal{B}^1(J_G, J_e)$ such that

$$\log |\mathcal{B}^{10}(J_G, J_e)| \geq c_1 J_e \log(eJ_G^2/J_e), \\ \rho(\mathbf{\Omega}_1, \mathbf{\Omega}_2) \geq c_2 J_e, \forall \mathbf{\Omega}_1, \mathbf{\Omega}_2 \in \mathcal{B}^{10}, \mathbf{\Omega}_1 \neq \mathbf{\Omega}_2$$

for some universal constants $c_1, c_2 > 0$. Then $\|\mathbf{\Omega}_1 - \mathbf{\Omega}_2\|_F^2 = \gamma^2 R^2 \rho(\mathbf{\Omega}_1, \mathbf{\Omega}_2) \geq c_2 \gamma^2 R^2 J_e$. It follows from the RIP that

$$\|\check{\mathbf{X}} \text{vec}(\mathbf{\Omega}_1) - \check{\mathbf{X}} \text{vec}(\mathbf{\Omega}_2)\|_2^2 \geq c_2 \underline{\kappa} \|\check{\mathbf{X}}\|_2^2 \gamma^2 R^2 J_e \quad (42)$$

for any $\mathbf{\Omega}_1, \mathbf{\Omega}_2 \in \mathcal{B}^{10}$, $\mathbf{\Omega}_1 \neq \mathbf{\Omega}_2$, where $\underline{\kappa}$ is a positive constant.

For Gaussian models, the Kullback-Leibler divergence of $\mathcal{N}(\check{\mathbf{X}} \text{vec}(\mathbf{\Omega}_2), \sigma^2 \mathbf{I})$ (denoted by $P_{\mathbf{\Omega}_2}$) from $\mathcal{N}(\check{\mathbf{X}} \text{vec}(\mathbf{\Omega}_1), \sigma^2 \mathbf{I})$ (denoted by $\mathcal{P}_{\mathbf{\Omega}_1}$) is $\mathcal{K}(\mathcal{P}_{\mathbf{\Omega}_1}, \mathcal{P}_{\mathbf{\Omega}_2}) = \frac{1}{2\sigma^2} \|\check{\mathbf{X}} \text{vec}(\mathbf{\Omega}_1) - \check{\mathbf{X}} \text{vec}(\mathbf{\Omega}_2)\|_2^2$. Let P_0 be $\mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$. By the RIP condition again, for any $\mathbf{\Omega} \in \mathcal{B}^1(J_G, J_e)$,

$$\mathcal{K}(P_0, P_{\mathbf{\Omega}}) \leq \frac{1}{2\sigma^2} \underline{\kappa} \|\check{\mathbf{X}}\|_2^2 \gamma^2 R^2 \rho(\mathbf{0}, \mathbf{\Omega}) \leq \frac{\gamma^2}{\sigma^2} \underline{\kappa} \|\check{\mathbf{X}}\|_2^2 R^2 J_e,$$

where we used $\rho(\mathbf{\Omega}_1, \mathbf{\Omega}_2) = J_e(\mathbf{\Phi}_1 - \mathbf{\Phi}_2) \leq 2J_e$. Therefore,

$$\frac{1}{|\mathcal{B}^{10}|} \sum_{\mathbf{\Omega} \in \mathcal{B}^{10}} \mathcal{K}(P_0, P_{\mathbf{\Omega}}) \leq \gamma^2 J_e \log(eJ_G^2/J_e) \leq \frac{\gamma^2}{c} \log\left(\frac{J_G^2}{J_e}\right). \quad (43)$$

Combining (42) and (43) and choosing a sufficiently small value for γ , we can apply Theorem 2.7 of Tsybakov (2009) to get the desired lower bound.

(ii) $J_G \log(ep/J_G) > J_e \log(eJ_G^2/J_e)$. Define a signal subclass

$$\mathcal{B}^2(J_G) = \{\mathbf{\Omega} = [\mathbf{b}, \mathbf{\Phi}^\top]^\top : b_j \in \{0, 1\} \forall j \in [p], \mathbf{\Phi} = \mathbf{0}, J_G(\mathbf{\Omega}) \leq J_G\},$$

where $R = \frac{\sigma}{\bar{\kappa}^{1/2} \|\bar{\mathbf{X}}\|_2} \{\log(ep/J_G)\}^{1/2}$ and $\gamma > 0$ is a small constant. The afterward treatment is similar to (i). The details are omitted.

For WH, we redefine \mathcal{B}^1 as $\mathcal{B}^1(J_G, J_e) = \{\mathbf{\Omega} = [\mathbf{b}, \mathbf{\Phi}^\top]^\top : b_j = 1 \text{ if } j \in [J_G] \text{ and } 0 \text{ otherwise, } \phi_{k,j} = 0 \text{ or } \gamma R \text{ if } k \geq j \text{ and } 1 \leq j \leq J_G, \text{ and } 0 \text{ otherwise, } J_e^w(\mathbf{\Phi}) \leq J_e\}$, where $R = \frac{\sigma}{\bar{\kappa}^{1/2} \|\bar{\mathbf{X}}\|_2} (\log(eJ_G p/J_e))^{1/2}$. Then $\mathcal{B}^1(J_G, J_e) \subset \text{WH}(J_G, J_e)$. The rest follows the same lines as the proof for SH.

B Coordinatewise error bound for support recovery

In this part, we show that GRESH estimators can recover the sparsity pattern of the true signal, i.e., $\mathcal{J}^* = \{j : \text{vec}(\mathbf{\Omega}^*)_j \neq 0\}$, with high probability. The result is of the type of Lounici (2008) and Ravikumar et al. (2010), but we do not have to assume the stringent irrerepresentable conditions or mutual coherence conditions. For more discussions, see Zhang (2009).

ASSUMPTION $\mathcal{C}(\mathcal{J}_e, \mathcal{J}_G, \zeta, \lambda_1, \lambda_2)$. Given $\mathcal{J}_e \subset [p^2]$, $\mathcal{J}_G \subset [p]$, $\zeta \geq 0$, $\lambda_1, \lambda_2 > 0$, for any $\mathbf{\Delta} = [\mathbf{\Delta}_b, \mathbf{\Delta}_\Phi^\top]^\top \in \mathbb{R}^{(p+1) \times p}$ satisfying $\mathbf{\Delta}_\Phi = \mathbf{\Delta}_\Phi^\top$, the following inequality holds

$$\begin{aligned} & \|\check{\mathbf{X}}\|_2 (\|(\mathbf{\Delta}_\Phi)_{\mathcal{J}_e}\|_1 + \|(\mathbf{\Delta}_{\mathcal{J}_G})^\top\|_{2,1}) + (J_e + J_G) \|\text{vec}(\mathbf{\Delta})\|_\infty^2 / (2\lambda_1 \wedge 2\lambda_2) \\ & \leq \|\check{\mathbf{X}}\|_2 (\|(\mathbf{\Delta}_\Phi)_{\mathcal{J}_e^c}\|_1 + \|(\mathbf{\Delta}_{\mathcal{J}_G^c})^\top\|_{2,1}) + \zeta \sqrt{J_e + J_G} \|\mathbf{X} \mathbf{\Delta}_b + \bar{\mathbf{X}} \text{vec}(\mathbf{\Delta}_\Phi)\|_2. \end{aligned} \quad (44)$$

Theorem 4. Assume $\boldsymbol{\epsilon} \sim N(\mathbf{0}, \sigma^2 \mathbf{I})$. Take $\lambda_1 = A_1 \sigma \sqrt{\log(ep)}$ and $\lambda_2 = A_2 \sigma \sqrt{\log(ep)}$ in (17) and suppose $(\mathbf{X}, \bar{\mathbf{X}}, \check{\mathbf{X}})$ satisfies $\mathcal{C}(\mathcal{J}_e^*, \mathcal{J}_G^*, \zeta, \lambda_1, \lambda_2)$ for

some $\zeta \geq 0$, where $\mathcal{J}_e^* = \mathcal{J}_e(\Phi^*)$, $\mathcal{J}_G^* = \mathcal{J}_G(\mathbf{b}^*, \Phi^*)$. Let $\hat{\Omega} = [\hat{\mathbf{b}}, \hat{\Phi}^\top]^\top$ be a global minimizer of (17). Define $\hat{\mathcal{J}} = \{j : |\text{vec}(\hat{\Omega})_j| > \sqrt{1 + \zeta^2}(\lambda_1 \wedge \lambda_2)\}$, and $\mathcal{J}^* = \{j : \text{vec}(\Omega^*)_j \neq 0\}$. Then for sufficiently large constants A_1 and A_2 , with probability at least $1 - Cp^{-c(A_1^2 \wedge A_2^2)}$, the estimate $\hat{\Omega}$ satisfies

$$\|\text{vec}(\hat{\Omega} - \Omega^*)\|_\infty \leq \sqrt{1 + \zeta^2}(\lambda_1 \wedge \lambda_2).$$

If, in addition, the minimum signal strength satisfies

$$\min_{j \in \mathcal{J}^*} |\text{vec}(\Omega^*)_j| > 2\sqrt{1 + \zeta^2}(\lambda_1 \wedge \lambda_2), \quad (45)$$

then $\hat{\mathcal{J}} = \mathcal{J}^*$, with probability at least $1 - Cp^{-c(A_1^2 \wedge A_2^2)}$.

To recover the true support, signal strength conditions like (45) must be imposed. Interestingly, compared with Theorem 5.1 in Lounici et al. (2011), hierarchical variable selection can accommodate smaller signals than group variable selection.

Proof. From the proof of Theorem 2, we get the following inequality with probability at least $1 - Cp^{-cL}$

$$\begin{aligned} & \frac{1}{2}M(\mathbf{b}^* - \hat{\mathbf{b}}, \Phi^* - \hat{\Phi}) + A\lambda^o\rho\|(\Delta_\Phi)_{\mathcal{J}_e}\|_1 + A\lambda^o\rho\|(\Delta_{\mathcal{J}_G^*})^\top\|_{2,1} \\ & \leq \frac{1}{2}\left(1 + \frac{3a_1}{A_0 - 6}\right)M(\mathbf{b}^* - \mathbf{b}, \Phi^* - \Phi) + \frac{1}{2}\left(\frac{1}{a_1} - 1\right)M(\Delta_b, \Delta_\Phi) \\ & \quad + A\lambda^o\rho\|(\Delta_\Phi)_{\mathcal{J}_e}\|_1 + A\lambda^o\rho\|(\Delta_{\mathcal{J}_G^*})^\top\|_{2,1} + \frac{a_1L}{2(1 - 6/A_0)}\sigma^2(J_e + J_G)\log(ep), \end{aligned}$$

where $\rho = \|\check{\mathbf{X}}\|_2$, L is a universal constant, $\lambda^o = \sigma\sqrt{\log(ep)}$, $A^2 = A_0L$, and a_1, A_0 are positive constants to be determined. Also recall that $\lambda_1 = A_1\lambda^o$, $\lambda_2 = A_2\lambda^o$, $A = A_1 \wedge A_2$.

By the regularity condition, we have

$$\begin{aligned}
& \frac{1}{2}(J_e + J_G) \|\text{vec}(\Delta)\|_\infty^2 + \frac{1}{2}M(\mathbf{b}^* - \hat{\mathbf{b}}, \Phi^* - \hat{\Phi}) \\
& \leq \frac{1}{2}(1 + \frac{3a_1}{A_0 - 6})M(\mathbf{b}^* - \mathbf{b}, \Phi^* - \Phi) + \frac{1}{2}(\frac{1}{a_1} - 1)M(\Delta_b, \Delta_\Phi) \\
& \quad + A\lambda^\circ \zeta \sqrt{J_e + J_G} \|\mathbf{X}\Delta_b + \bar{\mathbf{X}} \text{vec}(\Delta_\Phi)\|_2 + \frac{a_1 L}{2(1 - 6/A_0)} \sigma^2(J_e + J_G) \log(ep) \\
& \leq \frac{1}{2}(1 + \frac{3a_1}{A_0 - 6})M(\mathbf{b}^* - \mathbf{b}, \Phi^* - \Phi) + \frac{1}{2}(\frac{1}{a_1} + \frac{1}{a_2} - 1)M(\Delta_b, \Delta_\Phi) \\
& \quad + \left\{ \frac{a_1}{2(1 - 6/A_0)} + \frac{a_2 A_0}{2} \zeta^2 \right\} L \sigma^2(J_e + J_G) \log(ep),
\end{aligned}$$

for any $a_2 > 0$. Taking $\mathbf{b} = \mathbf{b}^*$, $\Phi = \Phi^*$ and choosing $a_1 = a_2 = 1$, $A_0 \geq 7$ give $(J_e^* + J_G^*) \|\text{vec}(\hat{\Omega} - \Omega^*)\|_\infty^2 \leq A_0(1 + \zeta^2) L \sigma^2(J_e^* + J_G^*) \log(ep)$. The conclusion follows. \square

C An ADMM algorithm

We describe an ADMM algorithm for solving problem (16). Recall $\bar{\mathbf{Z}} = [\mathbf{z}_1 \odot \mathbf{z}_1, \dots, \mathbf{z}_1 \odot \mathbf{z}_p, \dots, \mathbf{z}_p \odot \mathbf{z}_1, \dots, \mathbf{z}_p \odot \mathbf{z}_p]$, $\check{\mathbf{Z}} = [\mathbf{x}_1, \mathbf{z}_1 \odot \mathbf{z}_1, \dots, \mathbf{z}_1 \odot \mathbf{z}_p, \dots, \mathbf{x}_p, \mathbf{z}_p \odot \mathbf{z}_1, \dots, \mathbf{z}_p \odot \mathbf{z}_p]$, and $\Lambda_\Omega = \mathbf{1}\lambda_\Omega^\top$. To apply ADMM, we rewrite (16) into

$$\begin{aligned}
& \min_{\Omega, \Gamma, \Upsilon} \frac{1}{2} \|\mathbf{y} - \mathbf{X}\mathbf{b} - \text{diag}(\mathbf{Z}\Phi\mathbf{Z}^\top)\|_2^2 + \|\Lambda_\Omega^\top \odot \Upsilon^\top\|_{2,1} \\
& \quad + \{\|\lambda_b \odot \Gamma_b\|_1 + \|\Lambda_\Phi \odot \Gamma_\Phi\|_1\} \\
& \quad \text{s.t. } \Gamma_\Phi = \Gamma_\Phi^\top, \Gamma = \Omega, \Upsilon = \Omega,
\end{aligned} \tag{46}$$

where $\Upsilon \in \mathbb{R}^{n \times (p^2 + p)}$, $\Gamma = [\Gamma_b, \Gamma_\Phi^\top]^\top$ with $\Gamma_b \in \mathbb{R}^p$, $\Gamma_\Phi \in \mathbb{R}^{p^2 \times p^2}$. The augmented Lagrangian of (46) can be formed by use of two Lagrangian multiplier matrices ($\mathbf{L}_1, \mathbf{L}_2$) and a penalty parameter ρ :

$$\begin{aligned}
& L_\rho(\Omega, \Gamma, \Upsilon, \mathbf{L}_1, \mathbf{L}_2) \\
& = \frac{1}{2} \|\mathbf{y} - \check{\mathbf{Z}} \text{vec}(\Omega)\|_2^2 + \{\|\lambda_b \odot \Gamma_b\|_1 + \|\Lambda_\Phi \odot \Gamma_\Phi\|_1 + \iota_{\{\Gamma_\Phi = \Gamma_\Phi^\top\}}\} + \|\Lambda_\Omega^\top \odot \Upsilon^\top\|_{2,1} \\
& \quad + \rho \mathbf{1}^\top (\mathbf{L}_1 \odot (\Omega - \Gamma) + \mathbf{L}_2 \odot (\Omega - \Upsilon)) \mathbf{1} + \frac{\rho}{2} (\|\Omega - \Gamma\|_F^2 + \|\Omega - \Upsilon\|_F^2).
\end{aligned}$$

Based on the proof of Theorem 1, we can solve the sub-optimization problems of $\mathbf{\Gamma}$ and $\mathbf{\Upsilon}$ by some proximity operators; the details are omitted. The full ADMM algorithm is given as follows.

Algorithm 2 An ADMM algorithm

Inputs:

Data: \mathbf{X} , \mathbf{Z} , \mathbf{y} . Regularization parameters: $\lambda_b, \Lambda_\Phi, \lambda_\Omega$.

repeat

1. $\text{vec}(\mathbf{\Omega}) \leftarrow (\check{\mathbf{Z}}^\top \check{\mathbf{Z}} + 2\rho \mathbf{I})^{-1} [\check{\mathbf{Z}}^\top \mathbf{y} + \rho \text{vec}(\mathbf{\Gamma} + \mathbf{\Upsilon} - \mathbf{L}_1 - \mathbf{L}_2)]$;
2. $\mathbf{\Upsilon}[:, k] \leftarrow \Theta_S(\mathbf{\Omega}[:, k] + \mathbf{L}_2[:, k]; \lambda_\Omega[k]/\rho), \forall k : 1 \leq k \leq p$;
3. $\mathbf{\Omega}[2 : \text{end}, :] \leftarrow (\mathbf{\Omega}[2 : \text{end}, :] + \mathbf{\Omega}[2 : \text{end}, :]^\top)/2$;
4. $\mathbf{\Gamma} \leftarrow \Theta_S(\mathbf{\Omega} + \mathbf{L}_1; [\lambda_b, \Lambda_\Phi^\top]^\top/\rho)$;
5. $\mathbf{L}_2 \leftarrow \mathbf{L}_2 + \mathbf{\Omega} - \mathbf{\Upsilon}$;
6. $\mathbf{L}_1 \leftarrow \mathbf{L}_1 + \mathbf{\Omega} - \mathbf{\Gamma}$;

until convergence

Output: $\mathbf{\Omega}$

D Additional experiments

D.1 Microarray data

The dataset contains 21,575 genes collected from 298 mice (Wang et al., 2006). We use the body mass index (BMI) as the response variable. We chose 100 genes that are most correlated with BMI to study an interaction model, where $\mathbf{\Omega}$ contains 101×100 entries. The same hierarchial CV procedure is used as in the California housing data analysis. The corresponding results are shown in the following table.

Table 6: Performance of HL and GRESH on mouse BMI data

	Mean error	Median error	Median cardinality	Time
HL	12.8	9.06	51	4,138
GRESH	0.40	0.16	5	1,451

The top five main effects selected by GRESH more than 70% of the time correspond to FAM8A1, MLH1, VANG1, SEL1H and ASTL. However, except for

VANGL1, the other four almost never entered the HL models. The literature suggests that all these genes are in relation to BMI. For example, MLH1 mutations cause Lynch syndrome which typically results in a significantly lower BMI (Bruegl et al., 2014), and the insufficiency of SEL1H may predispose mice to high fat diet-induced hyperglycemia (Francisco1 et al., 2011). According to our experience in microarray data analysis, the influence of interactions is often not as large as that of main effects. When calling GRESH on the mouse data, the interaction between genes VANGL1 and RFWD2 occurred most often among all second-order terms, with however a low frequency of 10%.

References

- Bauschke, H. H. and P. L. Combettes (2008). “A Dykstra-like Algorithm for Two Monotone Operators”. *Pacific Journal of Optimization* 4(3), 383–391.
- Bickel, P. J., Y. Ritov, and A. B. Tsybakov (2009). “Simultaneous Analysis of Lasso and Dantzig Selector”. *The Annals of Statistics*, 1705–1732.
- Bickel, P. J., Y. Ritov, and A. B. Tsybakov (2010). “Hierarchical Selection of Variables in Sparse High-Dimensional Regression”. In *Borrowing Strength: Theory Powering Applications—A Festschrift for Lawrence D. Brown*, pp. 56–69. Institute of Mathematical Statistics.
- Bien, J., F. Bunea, and L. Xiao (2016). “Convex Banding of the Covariance Matrix”. *Journal of the American Statistical Association* 111(514), 834–845.
- Bien, J., J. Taylor, and R. Tibshirani (2013). “A Lasso for Hierarchical Interactions”. *The Annals of Statistics* 41(3), 1111–1141.
- Boyd, S., N. Parikh, E. Chu, B. Peleato, and J. Eckstein (2011). “Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers”. *Foundations and Trends® in Machine Learning* 3(1), 1–122.
- Bruegl, A., B. Djordjevic, D. Urbauer, S. Westin, P. Soliman, K. Lu, R. Luthra, and R. Broaddus (2014). “Utility of MLH1 Methylation Analysis in the Clinical Evaluation of Lynch Syndrome in Women With Endometrial Cancer”. *Current Pharmaceutical Design* 20(11), 1655–1663.

- Chipman, H. (1996). “Bayesian Variable Selection with Related Predictors”. *Canadian Journal of Statistics* 24(1), 17–36.
- Choi, N. H., W. Li, and J. Zhu (2010). “Variable Selection With the Strong Heredity Constraint and its Oracle Property”. *Journal of the American Statistical Association* 105(489), 354–364.
- Cohen, J., P. Cohen, S. G. West, and L. S. Aiken (2013). *Applied Multiple Regression/Correlation Analysis for the Behavioral Sciences*. Routledge.
- Davidson, E. H. and D. H. Erwin (2006). “Gene Regulatory Networks and the Evolution of Animal Body Plans”. *Science* 311(5762), 796–800.
- Francisco¹, A. B., R. Singh¹, H. Sha, X. Yan, L. Qi, X. Lei, and Q. Long (2011). “Haploid Insufficiency of Suppressor Enhancer Lin12 1-like (SEL1L) Protein Predisposes Mice to High Fat Diet-Induced Hyperglycemia”. *The Journal of Biological Chemistry* 286(25), 22275–22282.
- Hamada, M. and C. J. Wu (1992). “Analysis of Designed Experiments With Complex Aliasing”. *Journal of Quality Technology* 24(3), 130–137.
- Hao, N. and H. H. Zhang (2014). “Interaction Screening for Ultra-High Dimensional Data”. *Journal of the American Statistical Association* 109(1), 1285–1301.
- Hastie, T., R. Tibshirani, and J. Friedman (2009). *The Elements of Statistical Learning* (2nd ed.). New York: Springer-Verlag.
- He, Y., Y. She, and D. Wu (2013). “Stationary-Sparse Causality Network Learning”. *Journal of Machine Learning Research* 14(1), 3073–3104.
- Jenatton, R., J. Mairal, G. Obozinski, and F. Bach (2011). “Proximal Methods for Hierarchical Sparse Coding”. *Journal of Machine Learning Research* 12, 2297–2334.
- Laurent, B. and P. Massart (2000). “Adaptive Estimation of a Quadratic Functional by Model Selection”. *The Annals of Statistics*, 1302–1338.
- Lim, M. and T. Hastie (2015). “Learning Interactions via Hierarchical Group-Lasso Regularization”. *Journal of Computational and Graphical Statistics* 24(3), 627–654.

- Lounici, K. (2008). “Sup-Norm Convergence Rate and Sign Concentration Property of Lasso and Dantzig Estimators”. *Electronic Journal of Statistics* 2, 90–102.
- Lounici, K., M. Pontil, S. Van De Geer, and A. B. Tsybakov (2011). “Oracle Inequalities and Optimal Inference Under Group Sparsity”. *The Annals of Statistics* 39(4), 2164–2204.
- McCullagh, P. and J. A. Nelder (1989). *Generalized Linear Models*. London England Chapman and Hall.
- Negahban, S. N., P. Ravikumar, M. J. Wainwright, and B. Yu (2012). “A Unified Framework for High-Dimensional Analysis of M -Estimators With Decomposable Regularizers”. *Statistical Science* 27(4), 538–557.
- Nelder, J. A. (1977). “A Reformulation of Linear Models”. *Journal of the Royal Statistical Society. Series A (General)*, 48–77.
- Nesterov, Y. (2007). “Gradient Methods for Minimizing Composite Objective Function”. Technical report, Université catholique de Louvain, Center for Operations Research and Econometrics (CORE).
- Opial, Z. (1967). “Weak Convergence of the Sequence of Successive Approximations for Nonexpansive Mappings”. *Bulletin of the American Mathematical Society* 73(4), 591–597.
- Owen, A. B. (2007). “A Robust Hybrid of Lasso and Ridge Regression”. *Contemporary Mathematics* 443, 59–72.
- Pace, R. K. and R. Barry (1997). “Sparse Spatial Autoregressions”. *Statistics & Probability Letters* 33(3), 291–297.
- Peixoto, J. L. (1987). “Hierarchical Variable Selection in Polynomial Regression Models”. *The American Statistician* 41(4), 311–313.
- Peixoto, J. L. (1990). “A Property of Well-Formulated Polynomial Regression Models”. *The American Statistician* 44(1), 26–30.
- Radchenko, P. and G. M. James (2010). “Variable Selection Using Adaptive Nonlinear Interaction Structures in High Dimensions”. *Journal of the American Statistical Association* 105(492), 1541–1553.

- Ravikumar, P., M. J. Wainwright, and J. D. Lafferty (2010). “High-Dimensional Ising Model Selection Using ℓ_1 -Regularized Logistic Regression”. *The Annals of Statistics* 38(3), 1287–1319.
- Ravikumar, P. D., H. Liu, J. D. Lafferty, and L. A. Wasserman (2007). “SpAM: Sparse Additive Models”. In *In Advances in Neural Information Processing Systems 20*, pp. 1201–1208. MIT Press.
- Rigollet, P. and A. Tsybakov (2011). “Exponential Screening and Optimal Rates of Sparse Estimation”. *The Annals of Statistics* 39(2), 731–771.
- Schmidt, M., N. L. Roux, and F. R. Bach (2011). “Convergence Rates of Inexact Proximal-Gradient Methods for Convex Optimization”. In J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K. Weinberger (Eds.), *Advances in Neural Information Processing Systems 24*, pp. 1458–1466.
- She, Y. (2012). “An Iterative Algorithm for Fitting Nonconvex Penalized Generalized Linear Models With Grouped Predictors”. *Computational Statistics & Data Analysis* 56(10), 2976–2990.
- Simon, N., J. Friedman, T. Hastie, and R. Tibshirani (2013). “A Sparse-Group Lasso”. *Journal of Computational and Graphical Statistics* 22(2), 231–245.
- Tibshirani, R. (1996). “Regression Shrinkage and Selection via the Lasso”. *Journal of the Royal Statistical Society. Series B (Methodological)*, 267–288.
- Tsybakov, A. (2009). *Introduction to Nonparametric Estimation*. Springer Series in Statistics. New York: Springer.
- van de Geer, S. (2014). “Weakly Decomposable Regularization Penalties and Structured Sparsity”. *Scandinavian Journal of Statistics* 41(1), 72–86.
- Wang, S., N. Yehya, E. E. Schadt, H. Wang, T. A. Drake, and A. J. Lusis (2006). “Genetic and Genomic Analysis of a Fat Mass Trait With Complex Inheritance Reveals Marked Sex Specificity”. *PLoS Genet* 2(2), e15.
- Wei, F. and J. Huang (2010). “Consistent Group Selection in High-Dimensional Linear Regression”. *Bernoulli* 16(4), 1369–1384.

- Wu, J., B. Devlin, S. Ringquist, M. Trucco, and K. Roeder (2010). “Screen and Clean: A Tool for Identifying Interactions in Genome-Wide Association Studies”. *Genetic epidemiology* 34(3), 275–285.
- Yuan, M. and Y. Lin (2006). “Model Selection and Estimation in Regression With Grouped Variables”. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 68(1), 49–67.
- Zhang, C.-H. and J. Huang (2008). “The Sparsity and Bias of the Lasso Selection in High-Dimensional Linear Regression”. *The Annals of Statistics*, 1567–1594.
- Zhang, T. (2009). “Some Sharp Performance Bounds for Least Squares Regression With L_1 Regularization”. *The Annals of Statistics* 37(5A), 2109–2144.
- Zou, H. and T. Hastie (2005). “Regularization and Variable Selection via the Elastic Net”. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 67(2), 301–320.